

A Graphical View of Bayesian Variable Selection

Zaili Fang and Inyoung Kim*

March 8, 2013

Department of Statistics, Virginia Polytechnic Institute and State University, Blacksburg,
Virginia, U.S.A.

*To whom correspondence should be addressed:

Inyoung Kim, Ph.D.

Department of Statistics, Virginia Polytechnic Institute and State University, 410A Hutch-
eson Hall, Blacksburg, VA 24061-0439, U.S.A.

Tel: (540) 231-5366

Fax: (540) 231-3863

Email: inyoungk@vt.edu

Abstract

In recent years, Ising prior with the network information for the “in” or “out” binary random variable in Bayesian variable selections has received more and more attentions. In this paper, we discover that even without the informative prior a Bayesian variable selection problem itself can be considered as a complete graph and described by a Ising model with random interactions. There are many advantages of treating variable selection as a graphical model, such as it is easy to employ the single site updating as well as the cluster updating algorithm, suitable for problems with small sample size and larger variable number, easy to extend to nonparametric regression models and incorporate graphical prior information and so on. In a Bayesian variable selection Ising model the interactions are determined by the linear model coefficients, so we systematically study the performance of different scale normal mixture priors for the model coefficients by adopting the global-local shrinkage strategy. Our results prove that the best prior of the model coefficients in terms of variable selection should maintain substantial weight on small shrinkage instead of large shrinkage. We also discuss the connection between the tempering algorithms for Ising models and the global-local shrinkage approach, showing that the shrinkage parameter plays a tempering role. The methods are illustrated with simulated and real data.

Keywords: Cluster Algorithm; Global-Local Shrinkage; Graphical Model; Ising Model; KM Model; Long Tail Prior; Mixture Normals; Tempering Algorithm; Variable Selection.

Running Title :

1 Introduction

In this paper, we consider the standard multiple linear regression model $[\mathbf{y}|\boldsymbol{\beta}, \phi] \sim N(X\boldsymbol{\beta}, \phi^{-1}I)$, where \mathbf{y} is $n \times 1$ vector of the response variable, $X = (\mathbf{x}_1, \dots, \mathbf{x}_p)$ is an $n \times p$ matrix of predictors, $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)^T$ is $p \times 1$ model coefficient vector of the full model with $\beta_j, j = 1, \dots, p$ corresponding to the j th predictor, and ϕ is the precision parameter. The “in” or “out” of the predictor is represented by a binary random variable γ_j . The Bayesian spike and slab approaches to sampling γ_j ’s have been introduced by different authors and maintain one of the most active research areas in Bayesian statistics, such as the Stochastic Search Variable Selection (SSVS) (George and McCulloch, 1993) and rescaled spike and slab model (Ishwaran and Rao, 2005). In recent years, incorporating networked prior information of the predictor into those Bayesian variable selection models has received many attentions (Li and Zhang, 2010; Monni and Li, 2010; Stingo et al., 2011; Tai et al., 2010). In all these papers, the network information of the predictors are introduced through an informative prior for γ_j ’s, which is a binary random graph, but none of them treat the variable selection as a graphical model when the prior is noninformative. A binary random graphical model for the random vector $\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_p)^T$ is represented by an undirected graph $G = (V, E)$, where V represents the set of p vertices or nodes corresponding to p predictors and E is a set of edges connecting neighboring nodes. In this paper, based on a reparameterized Bayesian variable selection model, KM model (Kuo and Mallick, 1998), we generalize the Bayesian variable selection problem into a Bayesian graphical model, referred to Bayesian Variable Selection Graphical Model (BVGM), and demonstrate that with the noninformative prior for $\boldsymbol{\gamma}$ the model is essentially a complete graphical model.

The Markov chain random process on a random binary graph can be well modeled by a Ising model conditional on $\boldsymbol{\beta}$ and ϕ . Thus the posterior distribution of $\gamma_j = 1$ or the posterior distribution of j th predictor “in” the model can be achieved by sampling the random

binary variable. As one of the most active research areas, abundant theories and sampling procedures for Ising model have been reported. A nice review can be found in (Iba, 2001; Newman and Barkema, 1999). One difficulty to sample γ in BVGM is that the interactions are random since they are expressed by the product of β_j 's and ϕ . Another difficulty is due to the long-range interaction of the complete graph where each node is coupled or neighboring with all other nodes. In the literature, the well known approaches to handle random and long-range interactions are the cluster algorithm and a family of exchange Monte Carlo, parallel tempering and simulated tempering algorithm (Iba, 2001). For the issue of the cluster algorithm, Nott and Green (2004) introduced the Swendsen-Wang algorithm (Swendsen and Wang, 1987) into Bayesian variable selection and Monni and Li (2010) discussed the Wolff algorithm (Wolff, 1989) in the study of network-structured genomics data. However, both algorithms are constructed based on the graph prior for γ and consider fixed interaction only. Therefore, both are not applicable to the more general random complete graphical model. In this paper, we generalize the cluster algorithm based on Wolff's approach so that the cluster is formed even with the random interactions among nodes. Furthermore, our generalized Wolff algorithm is introduced for complete graph with noninformative prior for γ , and it is straightforward to combine the graphical prior information.

For the issue of tempering algorithm, so far to our best knowledge there are no work discussing the connection between the tempering algorithm and Bayesian variable selection. In all the Bayesian variable selection models, there is always a critical parameter associated with penalization or shrinkage. By showing the variable selection problem as a Ising model, we address that the well known shrinkage parameter in Bayesian variable selection is equivalent to the temperature parameter in a Ising model. However, in the regular tempering algorithm, there is only one global temperature as a random variable. In BVGM, we adopt the global shrinkage and local acting strategy (Polson and Scott, 2011, 2012), which is employed by

assigning the priors of scale normal mixtures for β_j 's (Barndorff-Nielsen et al., 1982; West, 1987). Each β_j has a local shrinkage parameter, its normal precision parameter, as the local temperature, and there is another global shrinkage parameter to place a constrain on all local parameters. Furthermore, assigning different prior for β_j 's precision parameter leads to different performance. The widely known priors for $p(\beta_j)$ in this area include Student-t (normal/gamma) prior (Tipping, 2001), Laplace (normal/inverse gamma) prior (Carlin and Polson, 1991; Hans, 2009; Park and Casella, 2008), horseshoe prior (Carvalho and Polson, 2010), and Jeffrey's prior (Bae and Mallick, 2004).

Another issue we concern in this paper is the dynamics of the selection probability under different shrinkage. In Bayesian variable selection with large p , instead of the appearance frequency of one of the 2^p possible models, usually the predictors are selected according to the posterior marginal selection probability, $p(\gamma_j = 1|\mathbf{y})$, since the frequency of one specific model is extremely small. We define the curves of the selection probabilities of all predictors against the shrinkage parameter as the profile curves of BVGM. These profile curves are important because they provide a direct view about how to select the shrinkage parameter. They also assess the performance of different priors for β . Unfortunately, we have not seen any work study the overall profile of selection probability under a wide range of shrinkage expect Lykou and Ntzoufras (2012) where they studied the selection probability against the shrinkage with Laplace prior only. Hence one purpose of this paper is to systematically study the dynamics of the selection probabilities, and compare different β_j priors with different weight on shrinkage. We address this issue by focusing on the orthogonal design. Interestingly, instead of priors with much weight on large shrinkage, our results indicates that the best performance of a prior is obtained by placing substantial weight on small shrinkage but not zero shrinkage. Among those β_j prior candidates, horseshoe prior is the one capable to maintain such shrinkage proportion for the widest range of shrinkage parameter, thus

considered as the best.

We also consider one extension of BVGM to Bayesian sparse additive model (BSAM). Unlike the popular topic of Bayesian variable selection, there are only few papers discuss Bayesian variable selection with nonparametric regression (Reich et al., 2009; Scheipl, 2011; Smith and Kohn, 1996). Based on the KM model, our BVGM is very straightforward to extended to BSAM. We employ the Lancaster and Šalkauskas (LS) spline basis (Chib and Greenberg, 2010; Lancaster and Šalkauskas, 1986) to express the nonparametric function components. To our best knowledge, our paper is the first one capable to connect the graphical model with the nonparametric regressors such that we can select an appropriate subset of the function components and estimate the flexible function curves simultaneously.

We first introduce the KM hierarchical model and full conditional distributions for sampling the parameters except γ in Section 2. In Section 3 we discuss the connection between Bayesian variable selection and binary random graphical model and express our model as the Ising model with noninformative prior for γ . Then in Section 4, we first introduce the single site algorithm for sampling γ , then present a generalized Wolff cluster algorithm. In Section 5, we focus on understanding the selection probability profile including the dynamics of the selection probability under different shrinkage priors, and we discuss the connection between the simulated tempering algorithm and priors of the scale mixture of normals. In Section 6 and 7 we consider two extensions, one is how to incorporate prior network information for γ , another one is how to extend to BSAM with LS basis. In Section 8 and 9, we illustrate our model with simulations and real data analysis. Finally, in the last section, we conclude our work and discuss other potential extensions of our model.

2 Bayesian Variable Selection with Normal Mixture Priors

We are interested in selecting a subset of predictors from the p potential candidates. Thus we introduce the binary random vector

$$\boldsymbol{\gamma} = (\gamma_1, \dots, \gamma_p)^T,$$

where $\gamma_j \in (0, 1), j = 1, \dots, p$ is the binary indicator random variable corresponding to the j th predictor. With $\gamma_j = 1$ we select predictor \mathbf{x}_j otherwise exclude it from the model. To implement the stochastic search for γ_j 's, SSVS considers a multi-mode point mass and Gaussian mixture prior for β_j 's, $[\beta_j | \gamma_j, \tau_\beta] \sim (1 - \gamma_j)\delta(0) + \gamma_j N(0, \tau_\beta^{-1})$, where $\delta(0)$ represents the point mass density at zero.

In this paper we consider the KM model, which is expressed as

$$\mathbf{y} = \sum_j^p \gamma_j \mathbf{x}_j \beta_j + \boldsymbol{\epsilon}, \quad (1)$$

where $\boldsymbol{\epsilon} \sim N(\mathbf{0}, \phi^{-1}I)$ is independent identical noise vector. We standardize the data set X and center the response \mathbf{y} such that $\sum_{i=1}^n x_{ij}^2 = 1, \sum_{i=1}^n x_{ij} = 0, j = 1, \dots, p$ and $\sum_{i=1}^n y_i = 0$. We may also include an intercept term μ in model (1) with a normal prior, which requires only a simple extra step in the sampling procedure. In Section 7, this parametric linear regression model is easy to extend to nonparametric additive model by using some basis function to express the j th individual function component with β_j as a parameter vector for the j th predictor.

The reasons we employ KM model in this paper are: first, it is more natural as a variable selection model, where $\gamma_j = 0$ indicates that j th predictor has no effect in the response. Second, spike and slab models such as SSVS consider a multi-mode prior for β_j 's which may have a mixing problem for sampling β_j 's since β_j 's may get trapped in the point mass mode

for a long time. This problem becomes worse when we extend the SSVS to nonparametric additive model (Scheipl, 2011), because the chance of moving between the point mass and the normal model for β_j becomes lower in higher dimensional space. The third reason can be demonstrated in next section where we can see that it is very straightforward to express a KM model in a Ising model, while it is difficult for SSVS models.

In usual Bayesian variable selection, the normal prior assigned to β_j 's has form $[\beta|\tau_\beta] \sim N(0, \tau_\beta^{-1} I_p)$, where τ_β is a common precision parameter for all β_j 's and usually assigned an gamma prior with scale $a/2$ and rate $b/2$. Similar prior can be $[\beta|\tau_\beta] \sim N(0, \tau_\beta^{-1}(X^T X)^{-1})$ assuming $\tau_\beta = g^{-1}\phi$, where g is a positive number called g -factor (Liang et al., 2008). Because of this simplicity, β_j 's and ϕ can be integrated out and a closed form of the posterior distribution for γ can be achieved. However, we realized this simplicity has many disadvantages for variable selection purpose. For example, if we integrate out τ_β and achieve the marginal prior of β as $p(\beta|a, b) \propto (\beta^T \beta + b)^{-\frac{p+a}{2}}$, we can see the prior of β_j 's are no longer statistically independent to each other. This is not a good idea to explore the whole joint distribution space of (γ, β) since the main purpose of Bayesian variable selection is to explore the space of γ , while dependent β prior will limit the stochastic searching space. Based on this argument, we follow the shrink globally act locally scheme suggested by Polson and Scott (2011) to assign independent normal mixture priors for β_j 's. In next section, we will see that the interactions of the Ising model are determined by β_j 's. The larger β_j 's of two predictors, the larger the interaction between them, then the corresponding nodes have high probability to be dependent, meaning they are either "aligned" (both γ_j equal to 1, or both equal to 0), or "anti-aligned" (one γ_j equals to 1, and another equals to 0). On the other hand, if β_j 's between two nodes are very small, then the two predictors are independent to flip their γ_j values. Therefore, we want β_j 's to be as flexible as possible to explore the configuration space, while we do not want to lose the control so we constrain the overall

variability of the interaction through a global parameter, which we refer to b .

The shrink globally act locally scheme is easy to be implemented by following hierarchical model with scale normal variance mixture priors for β_j 's.

$$\begin{aligned} [\beta_j | \tau_j, b] &\sim N(0, b^2 \tau_j^{-1}), \\ [\tau_j] &\sim p(\tau_j), \\ [\phi] &\sim p(\phi), \end{aligned} \tag{2}$$

where τ_j is the precision parameter for the conditional normal prior of β_j and plays the role of local tempering. $p(\tau_j)$ and $p(\phi)$ are the priors for τ_j 's and ϕ respectively. Similar hierarchical model in SSVS setting also has been discussed by Heaton and Scott (2010).

With these settings, we can easily achieve the full conditional distribution for $\beta_c \subseteq \beta$

$$[\beta_c | \mathbf{y}, \gamma, \beta_{\bar{c}}, \phi] \sim \begin{cases} N(\boldsymbol{\mu}_c, \Sigma_c) & \text{if } \gamma_c = \mathbf{1} \\ N(\mathbf{0}, D_c^{-1}) & \text{if } \gamma_c = \mathbf{0}. \end{cases} \tag{3}$$

Here we use a general subscript “ c ” to stands for subset of the index $\{1, \dots, p\}$. We use \bar{c} to present the complementary index set of c . In above expression, D_c is a $|c| \times |c|$ diagonal matrix with $\tau_j/b^2, j \in c$ as the diagonal elements, where $|c|$ stands for the cardinality of c . Σ_c and $\boldsymbol{\mu}_c$ are expressed as

$$\begin{aligned} \Sigma_c &= (\phi X_c^T X_c + D_c)^{-1}, \\ \boldsymbol{\mu}_c &= \phi \Sigma_c X_c^T (\mathbf{y} - X_{\gamma_{\bar{c}}} \beta_{\bar{c}}). \end{aligned} \tag{4}$$

With some notations abuse here, X_c stands for the sub-matrix of X corresponding the predictors in c , and $X_{\gamma_{\bar{c}}}$ is the sub-matrix of $X_{\gamma} = (\gamma_1 \mathbf{x}_1, \dots, \gamma_p \mathbf{x}_p)$ corresponding to predictors in \bar{c} .

We simply assign a noninformative prior for $\phi : [\phi] \sim \phi^{-1}$ and the posterior distribution of ϕ is simply a gamma distribution.

$$[\phi | \mathbf{y}, \beta, \gamma] \sim G\left(\frac{n}{2}, \frac{1}{2} \|\mathbf{y} - X_{\gamma} \beta\|^2\right). \tag{5}$$

The prior for τ_j 's are critical since it determines how the local action of the sampling process. Many different type of $p(\tau_j)$ can be considered. A very general review of different choice for $p(\tau_j)$ can be found in Polson and Scott (2010, 2011). In this paper we only consider three widely known $p(\tau_j)$'s that result in three typical marginal β_j 's priors with characteristics of heavy tail, heavy mass around zero and both. We refer to these marginal priors of β_j as Cauchy, Laplace and horseshoe priors which are achieved by assigning gamma prior $[\tau_j] \sim G(1/2, 1/2)$, inverse gamma prior $[\tau_j] \sim IG(1, 1/2)$ and half Cauchy prior $[\tau_j^{1/2}] \sim C^+(0, 1)$ to τ_j respectively. The density forms for these three normal mixture settings are list in Table 1 respectively. Notice, to avoid the confusion, the terms of ‘‘Cauchy’’, ‘‘Laplace’’ and ‘‘horseshoe’’ not only refer to the marginal priors of β_j 's but also represent the normal mixture settings. For example, in the context, ‘‘Cauchy prior’’ stands for normal/gamma setting such that the marginal prior of β_j is Cauchy and the prior for $p(\tau_j)$ is $G(1/2, 1/2)$.

Table 1: Summary of Cauchy, Laplace and horseshoe priors for the marginal prior of β_j 's, corresponding priors for $p(\tau_j)$ and the density functions of the shrinkage parameter κ_j .

Marginal prior	$p(\beta_j b)$	Prior for τ_j	Distribution for κ_j
Cauchy	$\pi b(\beta_j^2 + b^2)^{-1}$	$\tau_j^{-\frac{1}{2}} \exp(-\frac{\tau_j}{2})$	$\kappa_j^{-\frac{1}{2}} (1 - \kappa_j)^{-\frac{3}{2}} \exp(-\frac{b^2 \kappa_j}{2(1 - \kappa_j)})$
Laplace	$(2b)^{-1} \exp(- \beta_j /b)$	$\tau_j^{-2} \exp(-\frac{2}{\tau_j})$	$\kappa_j^{-2} \exp(-\frac{(1 - \kappa_j)}{2b^2 \kappa_j})$
Horseshoe	-	$\tau_j^{-\frac{1}{2}} (1 + \tau_j)^{-1}$	$\kappa_j^{-\frac{1}{2}} (1 - \kappa_j)^{-\frac{1}{2}} [1 - \kappa_j + b^2 \kappa_j]^{-1}$

By defining a scaleless parameter $b^* = b/\sqrt{\phi}$, The full conditional distribution for τ_j in Cauchy and Laplace settings are

$$[\tau_j|\beta_j, b] \propto \tau_j^{-3/2} \exp \left[-\frac{(\tau_j - b^*/|\beta_j|)^2}{2\tau_j b^{*2}/\beta_j^2} \right], \quad \text{Laplace prior} \quad (6)$$

$$[\tau_j|\beta_j, b] \propto \exp \left[-\frac{1}{2}(\beta_j^2/b^{*2} + 1)\tau_j \right], \quad \text{Cauchy prior} \quad (7)$$

and the full condition distribution for τ_j in horseshoe prior setting is obtained by

$$\begin{aligned}
[u_j|\beta_j, b, v_j] &\propto \exp \left[-\frac{1}{2} \left(\frac{\beta_j^2}{b^{*2}v_j} + 1 \right) u_j \right], \\
[v_j|\beta_j, b, u_j] &\propto v_j^{-1} \exp \left[-\frac{1}{2} \left(\frac{\beta_j^2 u_j}{b^{*2}} \frac{1}{v_j} + v_j \right) \right], \\
\tau_j &= u_j/v_j,
\end{aligned}
\tag{8}$$

Horseshoe prior

(6) is an inverse Gaussian distribution $ING(b^*/|\beta_j|, 1)$ with mean $b^*/|\beta_j|$ and shape parameter 1. (7) is an exponential distribution or Gamma distribution $G(1, (\beta_j^2/b^{*2} + 1)/2)$. The Gibbs sampler for horseshoe prior is implemented by using the redundant multiplicative reparameterization technique similar to Gelman (2006). Reparameterize τ_j as $\tau_j = u_j/v_j$ where u_j and v_j are independently distributed with prior $G(1/2, 1/2)$ respectively, then the prior for $\tau_j^{1/2} \sim C^+(0, 1)$, and the prior for β_j is the horseshoe prior. In (8), the full conditional distribution for u_j is a Gamma distribution $G(1, (\beta_j^2/(b^{*2}v_j) + 1)/2)$ and the full conditional distribution for v_j is generalized inverse Gaussian distribution $GING\left(0, \frac{\beta_j^2 u_j}{b^{*2}}, 1\right)$.

3 Bayesian Variable Selection and Binary Random Graphical Model

The noninformative prior for γ is $\gamma \sim (\frac{1}{2})^p$. Thus the full conditional distribution of $\gamma|\beta, \phi$ is directly derived from the likelihood of γ given β and ϕ . Given β , consider the matrix of marginal regression functions $R = (\mathbf{r}_1, \dots, \mathbf{r}_p) = (\beta_1 \mathbf{x}_1, \dots, \beta_p \mathbf{x}_p)$, with each column as the marginal regression vector for j th predictor vector. In additive nonparametric model (see Section 7), $\mathbf{r}_j = f_j(\mathbf{x}_j) = Z_j \beta_j$ is the nonparametric function component of \mathbf{x}_j expanding on the $n \times M_j$ basis matrix Z_j with $1 \times M_j$ coefficient vector β_j (M_j is the dimension of the basis). Here we consider parametric regression model only, thus the full conditional

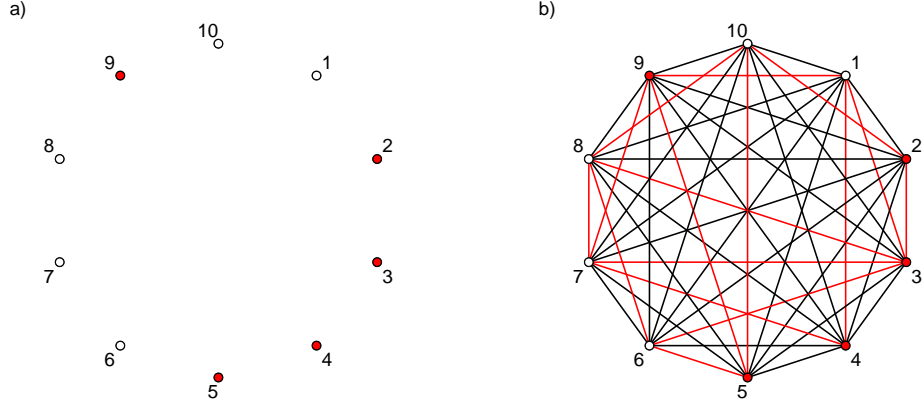


Figure 1: Diagram of variable selection as a random graph model with selected nodes (filled circles), excluded nodes (circles), edges of positive interaction (black lines), and edges of negative interaction (red lines). Independent variable selection: no interactions among nodes (a). General variable selection: a complete graph (b).

distribution of γ is

$$\begin{aligned}
 p(\gamma|\mathbf{y}, \beta, \phi) &\propto p(\mathbf{y}|\gamma, \beta, \phi) \\
 &\propto \exp\left(-\frac{1}{2}\phi\gamma^T R^T R\gamma + \phi\mathbf{y}^T R\gamma\right).
 \end{aligned} \tag{9}$$

This is nothing more than a Boltzman distribution of Ising model, $\frac{1}{Z}\exp(-U(\gamma))$, with

$$\begin{aligned}
 U(\gamma) &= -\gamma^T J\gamma - \mathbf{h}^T \gamma, \\
 J &= -\frac{\phi R^T R}{2}, \\
 \mathbf{h} &= \phi R^T \mathbf{y},
 \end{aligned} \tag{10}$$

where $Z = \sum_{\gamma} \exp(-U(\gamma))$ is called the partition (normalized) function and $U(\gamma)$ is called the energy of state γ given β and ϕ , J is the interaction matrix and \mathbf{h} is called “external field”. Above expression of Ising model is equivalent to following model:

$$p(\gamma|\mathbf{y}, \beta, \phi) \propto \exp\left(\sum_{i<j} J_{ij}\delta_{ij} + \sum_j h_j^* \gamma_j\right), \tag{11}$$

where the first summation is on all $i < j, j = 1, \dots, p$, $\delta_{ij} = 1$ if $\gamma_i = \gamma_j$ otherwise $\delta_{ij} = 0$, $J_{ij} = \phi \beta_i (\mathbf{x}_i^T \mathbf{x}_j) \beta_j$ is the non diagonal element of matrix J and h_j^* is the j th element of vector $\mathbf{h}^* = \phi R^T (\mathbf{y} - R\mathbf{1}/2)$. Above expression is achieved by plugging in following transformation into (9)

$$2 \left[\frac{1}{4} + \left(\gamma_i - \frac{1}{2} \right) \left(\gamma_j - \frac{1}{2} \right) \right] = \delta_{ij} = \begin{cases} 1 & \gamma_i = \gamma_j \\ 0 & \gamma_i \neq \gamma_j. \end{cases} \quad (12)$$

In the literature, the model with γ distributed as (9) is called spin glass model (consider γ_j has two spin states, up and down, corresponding to 1 and 0 respectively) when the coupling parameter J_{ij} follows some random distribution with positive or negative values. In our Bayesian variable selection model, because J_{ij} is the product of random variable β_j 's and ϕ each has a prior, the distribution for J_{ij} is some unknown distribution usually is neither iid nor tractable. Therefore, numerical method, such as MCMC sampling to simulate the distribution of γ is required. Now we can see that the choice of the prior for β_j 's is important since it directly effects the interaction among the nodes. The independent scale normal mixture prior for β_j 's is a nice choice since it is similar to the well known tempering algorithm in Ising model, and we can also derive some cluster algorithms. Both algorithms are expected to improve the mixing issue of the sampler (Nott and Green, 2004; Swendsen and Wang, 1987; Wolff, 1989).

Based on the Ising model, considering the p predictors as a set of nodes, we assign a binary random variable γ_j for each nodes. Those nodes may interact or couple with each other as described by a Ising model, so we have the following proposition:

Proposition 1: *The p dimension binary random variable $\gamma \in \{0, 1\}^p$ of the Bayesian variable selection problem based on KM model (1) is a class of stochastic processes on a finite random undirect graph model $G = (V, E)$, where $V = \{1, \dots, p\}$ is the set of nodes, corresponding to p predictors, and $E \subset V \times V$ is the set of edges. $\gamma \in \Gamma = \{(\gamma_1, \dots, \gamma_p) : \gamma_j \in (0, 1), j = 1, \dots, p\}$ is indexed by V with probability measure on Γ as (11), in which J_{ij} 's and*

h_j^* 's are all random with some distributions determined by the priori distributions of β_j 's and ϕ .

This is a complete graph model, since we don't limit the connection between any two nodes of V and the coupling between two nodes are long-range interaction. Figure 1 is the diagram of the graphical model for Bayesian variable selection. In Figure 1 (a), the interaction between any nodes $J_{ij} = 0$ thus this is a complete independent setting with which the configuration of γ depends on the "external field" \mathbf{h} only. Figure 1 (b) is a more general diagram for the complete graphical model. However, since any possible J is allowed, for a given J , a specific configuration of the edges will be given. For example, for a one dimension Ising model, the nodes form a one dimension chain, and one node only interacts with its two nearest neighbor nodes. This means the matrix J is a sparse matrix with non zero elements in positions $|i - j| \leq 1$ only. Furthermore, we can also consider the external field \mathbf{h}^* as a node indexed by 0 which represents the response variable \mathbf{y} , except that $\gamma_0 = 1$ is fixed. Then (11) can be expressed by a more compact form:

$$p(\gamma|\mathbf{y}, \beta, \phi) \propto \exp \left(\sum_{i < j} J_{ij} \delta_{ij} \right),$$

where $i, j = 0, 1, \dots, p$, and J_{ij} is the extended matrix with the first row and column equal to \mathbf{h}^* . However, in this paper we keep focus on expression (11) for explicitness.

4 Updating of γ

4.1 Single Site Algorithm

The joint posterior distribution of γ, β and ϕ will directly give this posterior distribution for γ . Because (9) and (11) are the direct form of Ising model, we can direct apply the Gibbs sampler procedure for γ based on (9) and (11) after sampling β and ϕ . This means we assign a noninformative prior for γ , $\gamma \sim \left(\frac{1}{2}\right)^p$, and the full conditional distribution given the data

for single site updating is

$$\begin{aligned}
[\gamma_j | \mathbf{y}, \boldsymbol{\gamma}_{\bar{j}}, \boldsymbol{\beta}, \phi] &\sim \text{Ber} \left(\frac{1}{1 + \pi} \right), \\
\pi &= \exp \left\{ -[U(\gamma_j = 1 | \boldsymbol{\gamma}_{\bar{j}}) - U(\gamma_j = 0 | \boldsymbol{\gamma}_{\bar{j}})] \right\} \\
&= \exp(-J_{jj} - 2J_{j\bar{j}}\gamma_{\bar{j}} - h_j),
\end{aligned} \tag{13}$$

where *Ber* stands for the Bernoulli distribution and $J_{j\bar{j}}$ is the j th row of J with j th column removed. $U(\gamma_j = 1 | \boldsymbol{\gamma}_{\bar{j}}) - U(\gamma_j = 0 | \boldsymbol{\gamma}_{\bar{j}})$ is the “energy” difference of two state configurations, $\gamma_j = 0 | \boldsymbol{\gamma}_{\bar{j}}$ and $\gamma_j = 1 | \boldsymbol{\gamma}_{\bar{j}}$, where $\boldsymbol{\gamma}_{\bar{j}}$ is the vector of $\boldsymbol{\gamma}$ with γ_j removed. Therefore the complete full conditional distributions of Gibbs sampler to update $\boldsymbol{\gamma}, \boldsymbol{\beta}$ and ϕ involves expression (3), (4), (5) and one of (6-8). This procedure is very simple and works well for most cases with moderate size p . The main advantage of our procedure is there is only one tuning parameter b , and the “tuning” process is extremely simple: just choose a b that separates the signals and noises with the largest gap in the marginal selection probability.

Above Gibbs sampler to update is one Metropolis-Hastings (MH) step with the Gibbs proposal and acceptance rate equal to one. We can also consider a MH one-step updating, which is more general in Ising model sampling. Denote the current state for γ_j as $\gamma_j^0 | \boldsymbol{\gamma}_{\bar{j}}$ and its flipped state $\gamma_j^* | \boldsymbol{\gamma}_{\bar{j}}$, whether or not we move from $\gamma_j^0 | \boldsymbol{\gamma}_{\bar{j}}$ to $\gamma_j^* | \boldsymbol{\gamma}_{\bar{j}}$ depends on the “energy” difference $\Delta U = U(\gamma_j^* | \boldsymbol{\gamma}_{\bar{j}}) - U(\gamma_j^0 | \boldsymbol{\gamma}_{\bar{j}})$. We prefer the system in lower “energy” state since the lower the energy the higher the probability. Thus if $\Delta U \leq 0$, the flipped state is accepted with probability 1. We treat the case $\Delta U > 0$ probabilistically, that is, with the probability to accept the flipped state as $p(\Delta U) = \exp(-\Delta U)$. These steps can be summarized as that we flip current state to its opposite rather than remaining the current state with probability

$$\min(1, \exp(-\Delta U)). \tag{14}$$

The detailed balance maintains and this MH updating is used in the MCMC Ising model sampling (Newman and Barkema, 1999; Nott and Green, 2004). In this paper, unless otherwise specified, we adopt this one step MH updating (14) with other Gibbs samplers in all




Figure 2: Diagram of the cluster algorithm. Forming the cluster (a-c). Flipping clustered nodes (d-e).

cases. Indeed, it is the antithetic updating method discussed by Nott and Green (2004) since $\exp(-\Delta U)$ is the odds of flipping current state with the Gibbs type proposal.

4.2 Cluster Algorithm

Beyond the single site algorithm, the cluster algorithm is well established for simulating model (9) when J_{ij} 's and h_j 's are fixed. There are abundant literatures available about the cluster algorithm in Ising model (Newman and Barkema, 1999; Swendsen and Wang, 1987; Wolff, 1989), but it has been introduced to Bayesian variable selection recently only

(Nott and Green, 2004). In general, a cluster algorithm performs better than the single site updating when J_{ij} is fixed. However, as pointed before, the model (9) is difficult in applying the clustering-updating algorithm since there is a random external field \mathbf{h}^* and the coupling coefficients J_{ij} 's follow some unknown distribution and are not independent. Plus the nodes are connected with each other by so called long-range interaction thus the system is a totally disordered complete graph. In this paper, we propose a generalized single-cluster Monte Carlo algorithm which is close to Wolff's clustering scheme but capable to handle the situation with long-range random interaction and external field.

In the original SW and Wolff algorithm, clusters are formed through the bonding between paired nodes with positive interactions. Although Nott and Green (2004) proposed an auxiliary variable technique to count the negative coupling between nodes and form clusters including anti-aligned nodes, their method is still based on the single bond between two nodes, which means whether adding a new node to the cluster is determined by the interaction between the new node and ONE node in the cluster. Unlike the usual Ising model on one dimension chain or two/three dimension lattice, the complete graph model of the binary random process is fully connected. This indicates each single node behaves according to the overall effects of all other nodes. Therefore, the clustering dynamics must incorporate this consideration. In other words, the growth of a cluster (adding one new node to the existing cluster) should consider the coupling between the new node and all nodes in the cluster.

Before introducing the cluster algorithm, we specify two types of clusters since the cluster is formed according to the coupling coefficient J_{ij} which can be either positive or negative.

- a cluster with nodes aligned.
- a cluster with nodes aligned and anti-aligned.

We use c to denote the cluster, and \bar{c} as the complement of c . The single node is considered

as special case of the cluster with aligned nodes. So within the second type of cluster, there are two sub clusters anti-aligned to each other. We denote these two sub clusters as c_1 and c_0 with $\gamma_{c_1} = \mathbf{1}$ and $\gamma_{c_0} = \mathbf{0}$ respectively.

The question then is, given a particularly defined probability p_a of adding a node to the cluster, what is the acceptance ratio that make the flip of the cluster satisfies detailed balance, and how to choose p_a such that the average acceptance ratio is as large as possible? So we derived following generalized Wolff algorithm based on these considerations.

1. Form the cluster.

(a) Initialize the cluster set c by randomly picking a seed node.

(b) Examine the nodes in \bar{c} one by one, add the node j in \bar{c} to the cluster with the probability

$$p_{a,j} = \max \left\{ 1 - \exp \left[\lambda (-1)^{\gamma_j} \left(\sum_{k \in c_1} J_{jk} - \sum_{l \in c_0} J_{jl} \right) \right], 0 \right\}, \quad (15)$$

and remove j from \bar{c} if j added to c , where $0 \leq \lambda \leq 1$. Continue iteratively until no new sites added when each nodes in \bar{c} has been examined.

2. Flip the nodes in cluster c with probability

$$\begin{aligned} & \alpha(\gamma_c^0 \rightarrow \gamma_c^*) \\ &= \min \left\{ \exp \left[(1 - \lambda) \sum_{j \in \bar{c}} (-1)^{\gamma_j} \left(\sum_{k \in c_1} J_{jk} - \sum_{l \in c_0} J_{jl} \right) + \sum_{j \in c_0} h_j^* - \sum_{j \in c_1} h_j^* \right], 1 \right\} \\ &= \min \left\{ \exp \left[(1 - \lambda) (\mathbf{1}^T J_{c_0 \bar{c}} - \mathbf{1}^T J_{c_1 \bar{c}}) (2\gamma_{\bar{c}} - \mathbf{1}) + \mathbf{1}^T \mathbf{h}_{c_0}^* - \mathbf{1}^T \mathbf{h}_{c_1}^* \right], 1 \right\}. \end{aligned} \quad (16)$$

3. Flip the rest nodes in \bar{c} (if any left) by single updating method (14).

4. Update β_j 's, τ_j 's and ϕ .

In (16), the last expression is for the convenience of coding using matrix expressions. As we can see, parameter λ plays a role of partial clustering similar to Higdon (1998). When

$\lambda = 1$, all interaction terms in (16) are annihilated, which means the coupling of the cluster with its neighbors are totally decoupled. If $\lambda = 0$, then no clustering process, the algorithm is reduced to single site algorithm.

The cluster algorithm can be better explained using the diagram in Figure 2. Figure 2 (a-c) demonstrate the clustering process. First we randomly select a seed node, in this diagram, node 8. Then we throw the bond to all neighbors of node 8, and find node 5 is bonded to 8 with probability $p_{a,5}$ and forms the cluster (the dashed line is turned into solid lines, meaning 5 is added to the cluster). We scan the remaining nodes again but whether or not a new node should be added is determined by the bonding between the new node and node 5 and 8. For example in Figure 2 (b), the bond between the new node 4 and the cluster is the overall bonds 4-5 and 4-8. In Figure 2 (c), after add the last new node 1 into the cluster, we scan all the left nodes and find no new node added to the cluster, then the clustering process stops.

The flipping of the cluster is demonstrated in Figure 2 (d-e). The cluster formed contains nodes $c = \{8, 5, 4, 2, 1\}$. To flip these nodes, we have to cut off the bonding of the cluster with all other nodes in \bar{c} because in a complete graph the neighbors of a cluster is all other nodes outside of the cluster. For example, the bond between the cluster and node 10 is demonstrated in Figure 2 (d), where we can see the bonds between 10 and all nodes in the cluster should be cut off to flip the cluster. Thus to completely flip the cluster, the bonds between all other nodes in \bar{c} and the nodes in c should be cut off. Similarly, in the reverse process to flip the cluster back, as shown in Figure 2 (e), all the bonds between the cluster c and the nodes in \bar{c} must be cut off.

It is easy to show that our algorithm is more general in sense that it is applicable to the complete graph with random interaction. When applied to Ising model on grid with positive fixed interaction J (only interactions among nearest neighbors account), our algorithm

evaluate to the original Wolff algorithm: the cluster growth by throwing bonds to nearest neighbors with probability $1 - \exp(-J)$.

Following theorem shows the algorithm stated above satisfies the detailed balance and ergodicity.

Theorem 1: *With the probability of adding node to the cluster, $p_{a,j}$, and the probability of moving from current configuration γ_c^0 to the flipped configuration γ_c^* , $\alpha(\gamma_c^0 \rightarrow \gamma_c^*)$, as defined as in the generalized Wolff algorithm, the algorithm is detailed balanced and ergodic.*

Proof: See A.1.

In this paper, we mainly focus on the noninformative prior for γ . However, since the distribution of γ given β and ϕ follows the Boltzman distribution, it is nature to assign a Boltzman prior or Ising prior for γ if such priori information is available. For example, in some genetic data, the genes form a network that can be descried using special graph model, with this information we can assign a Ising prior with specific interaction matrix that represents the priori graph structure. We will discuss this issue in Section 6. Another advantage of the cluster algorithm is it reveals the latent graph structure according to the frequencies of nodes that form a cluster, and this information may help us to distinguish the signals and the noise since the signals and noise should have high frequency to be anti-aligned.

5 Understanding the Mechanism of Bayesian Variable Selection

The purpose of this section is to understand how the marginal probability $p(\gamma_j|\mathbf{y})$ evaluates under different choice of marginal prior of β given the only tuning parameter b . Although our Ising model is based on the KM model (1), the results of this section is also applicable to SSVS model with the point mass mixture prior for β . This is because if we integrate

out β , both SSVS and KM models are identical. Note that all the results in this section is based on parametric linear model (1) where β_j is scalar, but the major results are similar to nonparametric linear model where β_j is $M_j \times 1$ vector.

Some notations are introduced here. Since \mathbf{x}_j 's are standarzed, $C = X^T X = [\mathbf{c}_1, \dots, \mathbf{c}_p]$ is the correlation matrix of \mathbf{x}_j 's and \mathbf{c}_j 's stands for the vector of the correlation between \mathbf{x}_j with all predictors. For orthogonal data set, $C = I_n$ or $\mathbf{x}_i^T \mathbf{x}_j = \delta_{ij}$ and $\mathbf{x}_j^T \boldsymbol{\epsilon} = 0, i, j = 1, \dots, p$. The projection of \mathbf{y} on \mathbf{x}_j can be expressed as $\mathbf{a} = X^T \mathbf{y} = (\mathbf{x}_1^T \mathbf{y}, \dots, \mathbf{x}_p^T \mathbf{y})^T = (a_1, \dots, a_p)^T$, which are the estimation of the signal β_j 's under orthogonal design. We may also need notation $\mathbf{c}_{\gamma_{\bar{j}}} = \gamma_{\bar{j}} \circ \{\mathbf{c}_j\}_{\bar{j}}$, where "o" stands for the pointwise product of two vectors, " \bar{j} " stands for the j th element removed for corresponding vectors and matrices,

5.1 General Profile of the Marginal Selection Probability

With the hierarchical model defined (1) and (2) in Section 2, the posterior distribution of β is multivariate normal given γ with mean $\boldsymbol{\mu}$ and variance Σ as.

$$\boldsymbol{\mu} = \phi \Sigma X_{\gamma}^T \mathbf{y}; \Sigma = (\phi X_{\gamma}^T X_{\gamma} + D)^{-1},$$

where D is a diagonal matrix with diagonal element $\{\frac{\tau_j}{b^2}\}_{1 \leq j \leq p}$. To understand how τ_j and b introduce the shrinkage effect, similar to Carvalho and Polson (2010); Polson and Scott (2011), it is convenient to introduce the shrinkage coefficient, $\kappa_j = \left(\frac{\tau_j}{b^2 \phi}\right) / \left(1 + \frac{\tau_j}{b^2 \phi}\right)$. Under the orthogonal design, the posterior mean and variance of β_j 's corresponding to $j \in \{j : \gamma_j = 1\}$ are

$$E(\beta_j | \mathbf{y}, \phi, b) = [1 - E(\kappa_j | \mathbf{y}, \phi, b)] a_j, \tag{17}$$

$$Var(\beta_j | \mathbf{y}, \phi, b) = [1 - E(\kappa_j | \mathbf{y}, \phi, b)] \phi^{-1}.$$

The coefficient κ_j 's represent how much shrinkage being placed on the initial estimation of β_j 's. $\kappa_j \rightarrow 0$, yields no shrinkage, and $\kappa_j \rightarrow 1$ yields near-total shrinkage. With this

definition of κ_j , it is easy to derive the density function of κ_j , $p(\kappa_j)$. Table 1 lists $p(\kappa_j)$'s based on the three prior settings given $\phi = 1$.

In order to compare the performance of different variance mixture priors $p(\tau_j)$ on the marginal selection probability and avoid notation abuses, it is convenient to assume ϕ fixed and use a scaleless transformation $\sqrt{\phi}\mathbf{y} \rightarrow \mathbf{y}$ such that $a_j\sqrt{\phi} \rightarrow a_j$, $b\sqrt{\phi} \rightarrow b$ and $\sqrt{\phi}C \rightarrow C$. This is equivalent to assume $\phi = 1$, but keep in mind that a_j 's, \mathbf{c}_j 's and b are scaled by $\sqrt{\phi}$ unless stated otherwise.




Figure 3: The curves of selection probability against κ_j (a). The curves of marginal selection probability against global shrinkage parameter b (b). Marginal selection probabilities with baseline subtracted (c). All plots are under orthogonal designs.

With these coefficients defined, following theorems connect the marginal odds of γ_j given the data with κ_j 's and b . Based on (1) and (2), the joint distribution for $\boldsymbol{\gamma}, \boldsymbol{\beta}, \boldsymbol{\tau}$ given b is

$$p(\boldsymbol{\gamma}, \boldsymbol{\beta}, \boldsymbol{\tau} | \mathbf{y}, b) \propto p(\mathbf{y} | \boldsymbol{\gamma}, \boldsymbol{\beta}) p(\boldsymbol{\beta} | \boldsymbol{\tau}, b) p(\boldsymbol{\tau}),$$

and the marginal probability for γ_j is $p(\gamma_j | \mathbf{y}, b) = \int \sum_{\boldsymbol{\gamma}_{\bar{j}}} p(\gamma_j, \boldsymbol{\gamma}_{\bar{j}}, \boldsymbol{\beta}, \boldsymbol{\tau} | \mathbf{y}, b) d\boldsymbol{\beta} d\boldsymbol{\tau}$. Thus the marginal odds for $\gamma_j = 1$ given b is

$$\pi_j^b = \frac{\int \sum_{\boldsymbol{\gamma}_{\bar{j}}} p(\gamma_j = 1, \boldsymbol{\gamma}_{\bar{j}}, \boldsymbol{\beta}, \boldsymbol{\tau} | \mathbf{y}, b) d\boldsymbol{\beta} d\boldsymbol{\tau}}{\int \sum_{\boldsymbol{\gamma}_{\bar{j}}} p(\gamma_j = 0, \boldsymbol{\gamma}_{\bar{j}}, \boldsymbol{\beta}, \boldsymbol{\tau} | \mathbf{y}, b) d\boldsymbol{\beta} d\boldsymbol{\tau}}.$$

Theorem 2: For the Bayesian model defined in (1) and (2), the marginal odds of γ_j , defined as π_j^b , has following form

$$\pi_j^b = \int \pi_j \xi_j p(\kappa_j) d\kappa_j, \quad (18)$$

where $p(\kappa_j)$ is the density function of κ_j ,

$$\pi_j = \kappa_j^{\frac{1}{2}} \exp \left[\frac{a_j^2}{2} (1 - \kappa_j) \right], \quad (19)$$

and ξ_j is a positive real function of κ_j

$$\xi_j = \frac{\int \sum_{\gamma_{\bar{j}}} \xi(\gamma_j = 1, \kappa_j, \gamma_{\bar{j}}, \tau_{\bar{j}}) p(\tau_{\bar{j}}) d\tau_{\bar{j}}}{\int \sum_{\gamma_{\bar{j}}} \xi(\gamma_j = 0, \kappa_j, \gamma_{\bar{j}}, \tau_{\bar{j}}) p(\tau_{\bar{j}}) d\tau_{\bar{j}}}. \quad (20)$$

1. For general cases

$$\xi(\gamma_j, \kappa_j, \gamma_{\bar{j}}, \tau_{\bar{j}}) = \exp \left[\frac{1}{2} \left(\mathbf{a}_{\bar{j}} - (1 - \kappa_j^{\gamma_j}) a_j \mathbf{c}_{\gamma_{\bar{j}}} \right)^T \Omega_j^{-1} \left(\mathbf{a}_{\bar{j}} - (1 - \kappa_j^{\gamma_j}) a_j \mathbf{c}_{\gamma_{\bar{j}}} \right) \right] |\Omega_j|^{1/2} |D_{\bar{j}}|^{1/2}$$

with $\Omega_j = [D_{\bar{j}} + X_{\gamma_{\bar{j}}}^T X_{\gamma_{\bar{j}}} - (1 - \kappa_j)^{\gamma_j} (\mathbf{c}_{\gamma_{\bar{j}}} \mathbf{c}_{\gamma_{\bar{j}}}^T)]^{-1}$.

For orthogonal designs, $\xi_j = 1$, and

$$\pi_j^b = \int \pi_j p(\kappa_j) d\kappa_j. \quad (21)$$

2. For orthogonal designs, if $\kappa_j \rightarrow 0$, then $\pi_j \rightarrow 0$, and if $\kappa_j \rightarrow 1$, then $\pi_j \rightarrow 1$. Similarly,

if $b \rightarrow 0$, then $\pi_j^b \rightarrow 1$, and if $b \rightarrow \infty$, then $\pi_j^b \rightarrow 0$.

Proof: See A.2.

From Theorem 2 we can see that in general the marginal odds $\pi_j^b \neq \int \pi_j p(\kappa_j) d\kappa_j$, the marginal odds of the orthogonal design. According to equation (18), when the correlation among predictors are not negligible, the odds will be “blurred” by the coefficient ξ_j , and the marginal selection probability is blurred too. Basically, ξ_j is a complex function of a_j , $\mathbf{c}_{\bar{j}}$, and $\tau_k/b^2, k \neq j$ or κ_j . Furthermore, it is infeasible to calculate ξ_j given large p with more than

2 predictors are correlated. However, we can focus on the orthogonal design to understand the mechanism of marginal selection probability in general since it is much more easier to calculate.

Combining Theorem 2 and Figure 3, we can understand the behaviors of π_j and π_j^b better. Figure 3 (a) plots the selection probability as a function of κ_j according to odds π_j (19), and Figure 3 (b) plots the marginal selection probabilities according to π_j^b with different prior $p(\kappa_j)$'s. We can see that for the orthogonal design, expression (19) and (21) indicate that π_j and π_j^b are monotone functions of a_j , this is demonstrated as the different selection probability curves in Figure 3 too. In ideal case, all noise predictors will demonstrate the same selection probability since $a_j = 0$, which defines the baseline selection probability curve in Figure 3 (a-b).

Thus ideally, any signals with $a_j \neq 0$ are deviated from the baseline curve. However, when correlations among variables do not equal to zero, the situations become complicated. First, even though the correlation among variables are small enough so π_j and π_j^b are still monotone function of a_j , the baseline will be blurred and extended to a band. To see this, consider $k \in \bar{V}^*$ and $j \in V^*$ where V^* is the set of true nodes and its complement is \bar{V}^* . Because $c_{kj} = \mathbf{x}_k^T \mathbf{x}_j \neq 0$, \mathbf{x}_k will have fake signal: $a_k = \mathbf{x}_k^T \mathbf{y} = c_{kj} a_j$. Thus all the noise predictors will demonstrate false signals as long as they have nonzero correlations with the true signals. This makes separating the true variable with small signals from the noise difficult. Secondly, because of ξ_j , even for large signals the selection probability will be distorted by their correlated fake signals. For example in Figure 3 (b), if $a_j = 2$ is the fake signal and $a_j = 4$ is the true signal and they are correlated, then the profile curve of $a_j = 2$ and $a_j = 4$ will show some “interacting” behavior at $b \approx 1$ where the selection probability of fake signal reaches the maximum (we will show this behavior in the simulation analysis). Thus in general the largest gap that separates $a_j = 2$ and $a_j = 4$ is not around $b \approx 1$, but in

two regions around $b \approx 0.1$ and $b \approx 1000$.

Furthermore, the second result of Theorem 2 states some asymptotic behaviors of π_j and π_j^b as $\kappa_j \rightarrow 0$ or $b \rightarrow \infty$ and $\kappa_j \rightarrow 1$ or $b \rightarrow 0$. This can be clearly seen in Figure 3 (a-b) where with small shrinkage ($\kappa_j \rightarrow 0$ or $b \rightarrow \infty$), both π_j and π_j^b approach 0, and with large shrinkage ($\kappa_j \rightarrow 1$ or $b \rightarrow 0$), they approach 0.5. However, the dropping rate depends on the magnitude of the signal and the prior $p(\kappa_j)$. For example in Figure 3 (a) we can see for large signal $a_j = 4$, the selection probability maintains at 1 for $\kappa_j \rightarrow 0$ till the last point. In Figure 3 (b), furthermore, we can see the selection probability curves are different for different priors: some drop very fast, such as Laplace prior, some are pretty robust to shrinkage such as horseshoe prior.

So choosing an appropriate prior for τ_j or $p(\kappa_j)$ is important. Our next question will be how to choose an appropriate prior. Based on Figure 3 and Theorem 2, there are some guidelines to choose the prior $p(\tau_j)$: (1) The rate of π_j^b to increase must be fast when the signal increases, so that the large true signal can be separated from the noise more easily. (2) π_j^b drops to 0.5 or 0 slowly when $b \rightarrow 0$ or $b \rightarrow \infty$ so we have a wider windows of b where the true signals maintain high selection probability.

Following theorems will further help us to understand the relationship between π_j^b and the shrinkage coefficient κ_j .

Theorem 3 For the Bayesian model (1) and (2) with orthogonal design, suppose prior $p(\beta_j)$ is a zero mean scale mixture of normals: $[\beta_j|\tau_j, b] \sim N(0, b^2\tau_j^{-1})$, with τ_j having proper prior $p(\tau_j)$. Define the marginal density $m_j = m(\mathbf{y}|\gamma_j = 1)$ as

$$m_j = \int p(\mathbf{y}|\beta_j, \gamma_j = 1)p(\beta_j|\tau_j)p(\tau_j)d\beta_jd\tau_j.$$

If m_j is finite for all \mathbf{y} , then

1.

$$E(\beta_j|\mathbf{y}, \gamma_j = 1) = a_j + \frac{1}{m_j} \mathbf{x}_j^T \frac{\partial}{\partial \mathbf{y}} m_j = a_j + \mathbf{x}_j^T \frac{\partial}{\partial \mathbf{y}} \log m_j. \quad (22)$$

2.

$$E(\beta_j|\mathbf{y}, \gamma_j = 1) = \frac{d}{da_j} \log \pi_j^b = [1 - E(\kappa_j|\mathbf{y}, \gamma_j = 1)]a_j. \quad (23)$$

Proof: See A.3

The first result of Theorem 3 is well known in Bayesian literature and can be found in Pericchi and Smith (1992) and Polson and Scott (2010) for more discussion, here we just simply extend it to the linear model case. We are more interested in the second result which gives the relationship among the expectation of β_j , the derivative of $\log \pi_j^b$ respect to a_j , and the shrinkage coefficient. Since we prefer a larger derivative of log marginal odds such that the large signal can be separate from the baseline further. (23) indicates that to achieve this purpose, it not only requires a large a_j , but also requires the expectation of shrinkage parameter κ_j to be small. This is confirmed by Figure 3 (a), where we see that the largest separation between the signals and the baseline is on the side of $\kappa_j \rightarrow 0$. Thus if integrate out κ_j to have π_j^b , we want the density $p(\kappa_j)$ has substantial mass around the region with largest separation. However, we don't want $\kappa_j = 0$ since it means exactly no shrinkage and all π_j 's drop to zero at this point.

Therefore, the general requirement for a $p(\kappa_j)$ based on Theorem 2, 3 and Figure 3 is to maintain substantial mass around region on the small shrinkage. Surprisingly, (23) seems contradict to the usual variable selection strategy that to recover the sparsity in the region of large shrinkage. In fact, it is possible to separate the signal and noise in large shrinkage region and large shrinkage does have some advantages, such as stability, faster mixing, less sensitive to nodes number p and so on. So it is a second choice as long as the signals are robust to large shrinkage, at least for large signals. However, in this paper we focus on the

small shrinkage region where the consistency in variable selection seems satisfied more often.

5.2 Dynamic Properties of the Odds with Different Priors

To explain the different behaviors of the selection probability caused by different priors, we need examine more details about the density distribution of $p(\kappa_j)$. Carvalho and Polson (2010) and Polson and Scott (2010) discussed the performance of different types of priors in the Bayesian regularization with difference weight on shrinkage. They focus on the effects on the estimation of the signals. We are looking at those priors from a different point of view in terms of variable selection based on the selection probabilities.

In the prior for β_j 's, b is the global parameter. As $b \rightarrow 0$, large global shrinkage is applied on all β_j 's, and as $b \rightarrow \infty$, the global shrinkage effect will be negligible. Table 1 lists the prior $p(\tau_j)$'s, and corresponding $p(\kappa_j)$'s as well as the marginal prior $p(\beta_j|b)$ s. Because of the existing of b , how much weight is put on the shrinkage is modified, and for different priors this modification is different.

To see this, Figure 4 compares density function $p(\beta_j|b)$'s around zero point and on the tails, and density function $p(\kappa_j)$'s given different b 's. By examining $p(\kappa_j)$'s in Figure 4 together with Figure 3 (c), we can understand how b effects the selection probability profile through putting different weight on shrinkage. Figure 3 (c) plots the selection probability profile with the baseline subtracted. It can be seen for orthogonal design, the larger the magnitude of the selection probability, the larger the true signal distinguished from the baseline. In small b ($b \leq 1$) region, the descendant order of the magnitude is Cauchy, horseshoe and Laplace prior for a given signal, which is consistent with the $p(\kappa_j)$ plots in Figure 4 at $b = 0.1$ where the order of density mass on the small κ_j region is Cauchy, horseshoe and Laplace prior. In addition, since Cauchy and horseshoe priors put similar mass around small κ_j side, their selection probabilities behave almost identically for small




Figure 4: Marginal prior density functions of β_j and density functions of κ_j for different b .

b as shown in Figure 3 (c). On the other hand, this order changes for $b = 100$ where it becomes horseshoe, Laplace and Cauchy prior in Figure 4. Again this is consistent with the selection probability order in Figure 3 (c) for large b ($b \geq 100$). The reason that Cauchy prior becomes worse for large b is because all the mass of $p(\kappa_j)$ is absorbed to $\kappa_j = 0$ which is not we expect as mentioned before. For a moderate b , such as $b = 1$, all priors have substantial mass around small κ_j side as shown in Figure 4, thus all behave similarly. This is confirmed by Figure 3 (c) where the selection probabilities for different prior seems similar around $b = 1$ at least for large signals.

Above analysis also shows that more weight on large shrinkage is not as important as on small shrinkage in terms of distinguishing the signals. Therefore, horseshoe prior is superior to other two, even though $p(\beta_j|b)$ of horseshoe prior does not have long tail as much as Cauchy prior for small b . Horseshoe prior does demonstrate that for a wide range of b $p(\kappa_j)$ maintains substantial mass on the small shrinkage side of κ_j . On the other hand, Laplace prior has almost zero mass around small κ_j side when b is small, and Cauchy prior has all mass absorbed to $\kappa_j = 0$ when b is very large, each deteriorates their performance for those b values respectively. Our argument to evaluate the priors is thus different from Carvalho and Polson (2010) where they argue that the horseshoe prior is superior because it has substantial mass on both small shrinkage and large shrinkage in terms of estimation. Of course, although we prefer small shrinkage in terms of variable selection, large shrinkage does have advantages that some times we must consider. For example, we found in the simulation that with large shrinkage the Gibbs sampler can converge faster even with very large p .

To further examine the dynamics of the selection probability profile, following theorem gives some asymptotic behaviors about the derivation of $\log \pi_j^b$ respect to $|a_j|$ and b , and so it helps us evaluate different priors.

Theorem 4: Consider the inverse of τ_j , $\sigma_j^2 = \tau_j^{-1}$, has prior density, $p(\sigma_j^2)$, as $\sigma_j^2 \rightarrow \infty$,

$$p(\sigma_j^2) \sim (\sigma_j^2)^{\alpha-1} e^{-\lambda \sigma_j^2} L(\sigma_j^2) d\sigma_j^2$$

for some slowly varying function $L(x)$ such that as $x \rightarrow \infty$ for all $t > 0$, $L(tx)/L(x) \rightarrow 1$, then

1. as $a_j \rightarrow \infty$

$$\frac{d}{d|a_j|} \log \pi_j^b \sim \begin{cases} |a_j| + \frac{2\alpha-1}{|a_j|} & \text{if } \lambda = 0 \\ |a_j| + \frac{\alpha-1}{|a_j|} - \frac{\sqrt{2\lambda}}{b} & \text{if } \lambda > 0. \end{cases} \quad (24)$$

2. For large a_j , as $b \rightarrow 0$

$$\frac{d}{db} \log \pi_j^b \sim \begin{cases} \frac{d}{db} \log L^b(a_j^2) & \text{if } \lambda = 0 \\ \frac{\sqrt{2\lambda}}{b^2} |a_j| + \frac{d}{db} \log L^b(|a_j|) & \text{if } \lambda > 0. \end{cases} \quad (25)$$

where $L^b(x)$ is a function conditioning on b (see A.4)

Proof: See A.4

Particulary, a $L^b(x)$ has forms of $b \exp\left(-\frac{b^2}{2x}\right)$, b^{-2} , and $\frac{b}{b^2+x^2}$ for Cauchy, Laplace and horseshoe prior respectively. When $|a_j|$ is large, as Theorem 4 assumes, $\frac{d}{db} \log L^b$ is negligible.

Theorem 4 is similar to the tail robustness theorem discussed by Polson and Scott (2011) about marginal density $m(\mathbf{y}|\gamma_j = 1)$, which implies that the shrinkage will vanish for any scale mixture normals with $p(\sigma_j^2)$ with heavier tails (such as Cauchy and horseshoe prior), while remain non-diminishing for $p(\sigma_j^2)$ with exponential tails (such as Laplace prior). Combining with (22) of Theorem 3, we get the similar conclusion about the estimation of β_j that it is robust if estimated by long tail priors. Similar robustness can be found for π_j^b . The robustness of π_j^b means fast change rate of π_j^b as signal magnitude increases and small change rate of π_j^b as shrinkage increases, which are important since these two characteristics can make distinguishing the signals easier. Large $\frac{d}{d|a_j|} \log \pi_j^b$ helps to distinguish the signals from the baseline, while small $\frac{d}{db} \log \pi_j^b$ leads to a wide window of b where the selection probability of true signals remain highly.

Expression (24) indicates that for priors with exponential tails ($\lambda > 0$), the selection probability π_j^b increase with a smaller rate when the signal magnitude increases comparing with the heavier tail priors. Meanwhile, expression (25) shows that for priors with exponential tails, π_j^b drops with much faster rate ($\sim \sqrt{2\lambda}/b^2$) as $b \rightarrow 0$. We can also compare the dropping rates of $\log \pi_j^b$ of Cauchy and horseshoe prior as $b \rightarrow 0$. Since $L^b(x) = \exp\left(-\frac{b^2}{2x^2}\right) b$ and $(1+b^2/x^2)^{-1}b$ for Cauchy and horseshoe prior respectively, it turns out $d \log L^b(x^2)/db \approx 1/b$ as $b \rightarrow 0$ for both priors. This means the dropping rate as $b \rightarrow 0$ is similar for Cauchy prior

and horseshoe prior which is confirmed in Figure 3 (b). Note the second conclusion of Theorem 4 does not apply to $b \rightarrow \infty$ unless $|a_j| \rightarrow \infty$ faster than $b \rightarrow \infty$, i.e., $b/|a_j| \rightarrow 0$ so L^b maintains as slowly varying function. However, as shown by the exactly calculation in Figure 3 (b), horseshoe prior is also the most robust one as $b \rightarrow \infty$.

Figure 5 gives the exact calculation of $r_a = \frac{d}{d|a_j|} \log \pi_j^b$ and $r_b = \frac{d}{db} \log \pi_j^b$ for three priors. In Figure 5 (a) we can see that r_a is the nearly the same for three priors at large b . However, when b is small, r_a is reduced by certain value for Laplace prior, meanwhile, it remains the same for Cauchy and horseshoe prior. So the exact calculation just confirms Theorem 4. Similarly, the exact calculation also confirms the result of Theorem 4 about r_b . As we can see in Figure 5 (b), r_b increases exponentially as $b \rightarrow 0$ for Laplace prior, which means as $b \rightarrow 0$, the selection probability by Laplace prior will exponentially drop to 0.5, and this behavior has already been observed in Figure 3 (b).

Based on discussion in Section 5.1 and this section, horseshoe prior performs the best in terms of the marginal selection probability, Cauchy prior is in the second place, and Laplace prior is the worst since the selection probability drops too fast as shrinkage increases.

5.3 Some Expressions for π_j^b

In above sections, we discussed the properties of the odds π_j^b for orthogonal designs, but did not show how to calculate π_j^b . Those curves are calculated by Monte Carlo simulations which are very precise. In some cases, we may want to calculate π_j^b directly. There is no closed form of π_j^b for the three different priors. However, we can see at least for Laplace and horseshoe priors, π_j^b can be expressed by some special functions.

For Laplace prior,

$$\pi_j^b = \sqrt{\frac{\pi}{2}} b^{-1} \exp \left[\frac{(|a_j| - b^{-1})^2}{2} \right] \left\{ \Phi(|a_j| - b^{-1}) + \exp(2|a_j|b^{-1}) \Phi(-|a_j| - b^{-1}) \right\}, \quad (26)$$

where $\Phi(\cdot)$ is the CDF of standardized normal distribution. This expression can be directly

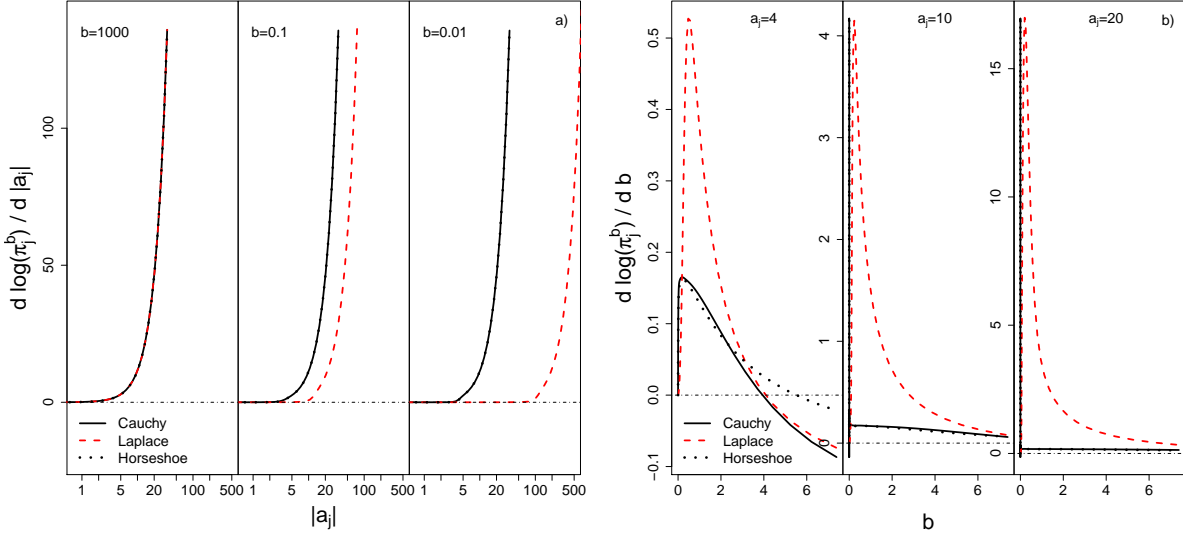


Figure 5: The derivative of log odds respect to $|a_j|$ against given different b (a). Note the curves of Cauchy and horseshoe priors are overlapped. The derivative of log odds respect to b given different a_j (b).

used to calculated the marginal selection probability given b .

For horseshoe prior,

$$\pi_j^b = \frac{1}{\pi} b^{-1} Be\left(1, \frac{1}{2}\right) \Phi_1\left(\frac{1}{2}, 1, \frac{3}{2}, \frac{a_j^2}{2}, 1 - b^{-2}\right), \quad (27)$$

where $Be(\dots)$ denotes the beta function, and $\Phi_1(\dots)$ is the degenerate hypergeometric function of two variables (Gordy, 1998; Polson and Scott, 2010). The calculation of Φ_1 can be employed by using a series of hypergeometric ${}_2F_1$ functions (Gordy, 1998).

The derivative of above expressions is shown in A.5. π_j^b of Cauchy prior does not have an analytic form, and its can not be represented by known special functions neither. Hence we simply use the Monte Carlo approach to calculate π_j^b for Cauchy prior.

5.4 Simulated Tempering and Generalization by Lévy Process

Li and Zhang (2010) discussed the difficulty of sampling around phase transition in a SSVS

model by assigning a Ising prior for γ . The difficulty is, given a Ising model there is a threshold for the interaction strength, when the interaction magnitude is larger than this threshold, the MCMC sampling will dramatically slow down, resulting in either overwhelming many selected nodes or extremely few ones. It becomes even worse when J_{ij} 's and h_j 's are all random, such as our model. However, the family of exchange Monte Carlo and simulated tempering algorithm has been developed to handle the slow mixing problem (Geyer and Thompson, 1995; Iba, 2001; Lyubartsev et al., 1992). By introducing the scale normal mixture for β , our model is an special simulated tempering algorithm which thus improves the mixing issue too.

To understand the simulated tempering algorithm, consider the usual Ising model with $U(\gamma, J) = \sum_{i < j} J_{ij} \delta_{ij}$ (for simplicity no external field h_i^* included), then the Boltzman distribution is expressed as

$$p(\gamma|T, J) = \frac{1}{Z(T)} \exp \left[-T^{-1} U(\gamma, J) \right],$$

where T represents the temperature (or the scale of variation), and J_{ij} is random and follows some distribution $p(J_{ij})$ such as standard Gaussian distribution. When $T \rightarrow 0$, the effective interaction $\tilde{J}_{ij} = T^{-1} J_{ij} \rightarrow \infty$. Thus if T is lower than some critical temperature, the strong interaction will lead to some non-ergodic behavior such as the slow down of the MCMC and extremely large proportion of $\gamma_j = 1$. The reason for this is because the low temperature phase of disordered Ising model generally has numerous local minima which are separated to each other by energy barriers. The characteristic time in which the system escapes from a local minimum, however, increases rapidly as the temperature decreases or the interaction increases. A good review can be found at Newman and Barkema (1999) about this issue. The family of tempering algorithm treats temperature T as a dynamical variable (Lyubartsev

et al., 1992), and the joint distribution $p(\boldsymbol{\gamma}, T)$ is represented as

$$p(\boldsymbol{\gamma}, T, J) \propto p(\boldsymbol{\gamma}|T, J) \prod_{i < j} p(J_{ij}) p(T), \quad (28)$$

where $p(T)$ is the distribution of T . The prior information for the T thus represents the range and mass of the temperature to sample the MCMC. With some variable transformation by replace $T_b^{-1/2} J_{ij}$ with \tilde{J}_{ij} , where $T^2 = T_b$, the joint distribution (28) then becomes

$$p(\boldsymbol{\gamma}, T_b, \tilde{J}) \propto p(\boldsymbol{\gamma}|\tilde{J}) \prod_{i < j} p(\tilde{J}_{ij}|T_b) T_b^{-1/2} p(T_b),$$

where $p(\boldsymbol{\gamma}|\tilde{J}) \propto \exp \left[- \sum_{ij} \tilde{J}_{ij} \delta_{ij} \right]$, $p(\tilde{J}_{ij}|T_b) \propto p(T_b^{1/2} \tilde{J}_{ij}) T_b^{1/2}$ (with some notation abuse, the later $p(\cdot)$ represents the same density function of $p(J_{ij})$). Clearly, T_b is a global temperature parameter here. If we introduce the local temperature parameter T_b for each interaction J_{ij} , then the marginal prior for \tilde{J}_{ij} is

$$p(\tilde{J}_{ij}) \propto \int_0^\infty p(\tilde{J}_{ij}|T_b) T_b^{-1/2} p(T_b) dT_b.$$

If $p(J_{ij}) \sim N(0, 1)$, above posterior for \tilde{J}_{ij} is a normal scale-mixture whose mixing measure is expressible in terms of the density of the subordinator T_b . Hence according to the Theorem 3 of Polson and Scott (2011), with the simulated tempering algorithm the interaction of the random Ising model (28) can be expressed as a Lévy process mixture scaled by T_b , and T_b is a nondecreasing pure-jump Lévy process with marginal density $p(T_b)$ at time b .

As another algorithm in the same family, the exchange monte carlo algorithms (or parallel tempering) is to simultaneously and independently simulate $K \geq 2$ replicas of the MCMC trace under different temperatures, and exchange the $\boldsymbol{\gamma}$ configurations of the replicas with certain acceptance probability by referring to the energy cost ΔU . The analogy between the exchanged monte carlo and simulated tempering algorithm is clear in terms of the mixture distribution of the interaction J_{ij} . For simulated tempering the mixture weight is the continuous prior $p(T_b)$ while the exchanged monte carlo is mixed with weight on a set of discrete

temperatures. In both algorithms, the low temperature process can access a representative set of local energy minimums with the accompany of the high temperature process which are generally able to sample large volumes of configuration space to keep the configuration from trapping in some local minimum.

We see that how to understand the simulated tempering algorithm as the Ising model with normal scale-mixture prior mixed by the Lévy process. On the other hand, BVGM with Lévy process mixtures can also be understood as an Ising model sampled by simulated tempering algorithm. To see this, we can generalize both Cauchy and Laplace prior into the framework of normal/generalized inverse Gaussian mixture. The marginal prior of β_j for both priors can be expressed using one formula

$$\begin{aligned}
p(\beta_j|u, v, w) &= \int_0^\infty p(\beta_j|\tau_j) \tau_j^{-1} g(\tau_j) d\tau_j = \int_0^\infty p(\beta_j|\tau_j) p(\tau_j|u, v, w) d\tau_j \\
&= \int_0^\infty \frac{1}{\sqrt{2\pi}} \exp(-\beta_j^2 \tau_j / 2) \tau_j^{1/2} \frac{(u/v)^w}{2K_w(uv)} \tau_j^{w-1} \exp[-(u^2 \tau_j + v^2 \tau_j^{-1}) / 2] d\tau_j \\
&= \frac{1}{\sqrt{2\pi}} \cdot \frac{(u/v)^w}{K_w(uv)} \cdot \frac{K_{w+1/2}\left(v\sqrt{\beta_j^2 + u^2}\right)}{\left(\sqrt{\beta_j^2 + u^2/v}\right)^{w+1/2}},
\end{aligned} \tag{29}$$

where $K_w(x)$ denotes the modified Bessel function of the third kind with w . $p(\tau_j|u, v, w)$ is the generalized inverse Gaussian ($GING(w, u, v)$) distribution with parameters u, v and w such that $w \in \mathbb{R}$ while u and v are both nonnegative and not simultaneously 0. Note for the two special cases we adopted in this paper, Cauchy and Laplace prior, the values for these there parameters are on the boundary. However, it turns out both the prior $p(\tau_j|u, v, w)$ and the marginal prior $p(\beta_j|u, v, w)$ as the limit exist. For Cauchy prior, $w = 1/2$, $u = b$, and $v = 0$, thus $p(\tau_j|b, v, 1/2) \rightarrow G(1/2, b^2/2)$ as $v \rightarrow 0$, where we use identity $K_{-1/2}(x) = \sqrt{\pi/2} x^{-1/2} \exp(-x)$. Hence the distribution of β_j reduces to the Cauchy distribution with scale parameter b , where we use $\lim_{x \rightarrow 0} x^w K_w(x) \rightarrow 2^{w-1} \Gamma(w)$ and $\Gamma(w)$ is a Gamma function. For Laplace prior, $w = -1$, $u = 0$, and $v = b^{-1}$. the limit

of prior $p(\tau_j|u, b^{-1}, -1) \rightarrow IG[1, (2b^2)^{-1}]$ as $u \rightarrow 0$, where $IG(\cdot)$ stands for inverse gamma distribution, and we used the index symmetry $K_{-w}(x) = K_w(x)$.

According to Theorem 3 of Polson and Scott (2011), we see $\tau_j^{-1}g(\tau_j) = p(\tau_j|u, v, w)$ where $g(\tau_j)$ is the density of the subordinator τ_j at (u, v, w) . Analogous to the discussion with the simulated tempering algorithm for Ising model, we can see the temperature parameter in our model is τ_j . However, there are several differences, such as in our model, $J_{ij} \propto \beta_i\beta_j$ and we assign normal mixture prior for β_j 's, while in the simulated tempering algorithm of regular Ising model, the prior is assigned to J_{ij} directly.

Now we can understand the temperature effect of τ_j . When $\tau_j \rightarrow \infty$, which is equivalent to $\kappa_j \rightarrow 1$, the system is in a high temperature state. This means MCMC is exploring the whole configuration space, and no precise sampling for a local energy minimum. This is also equivalent to say for each node, the odds of $\gamma_j = 1$ is equal to one since in high temperature every node is heated up and chance to be up and down is even, which means the marginal selection probability for all nodes is $1/2$ as shown in Figure 3 (a) as $\kappa_j \rightarrow 1$. On the other hand, when $\tau_j \rightarrow 0$, which is equivalent to $\kappa_j \rightarrow 0$, the system is in low temperature state. Starting from some initial state, all nodes configure will be trapped into their energy minimum (local maximum likelihood) which is $\gamma_j = 0$ for most nodes in the orthogonal design, unless the external field $h_j \propto a_j$ is strong enough to force the node in the state $\gamma_j = 1$. This is why we see in Figure 3 (a) for small κ_j the selection probability is 0 for nodes with small a_j and remains 1 for nodes with very large a_j .

It is easy to understand the role of b too, which is opposite to τ_j if we look at the prior $p(\beta_j|\tau_j, b) \sim N(0, b^2/\tau_j)$. In fact, we can also understand b by representing the hierarchical model as $p(\boldsymbol{\gamma}|\mathbf{y}, \boldsymbol{\beta}) \prod p(\beta_j|\tau_j)p(\tau_j|b)$, where $p(\tau_j|b) = \tau_j^{-1/2} \exp(-b^2\tau_j/2)$, $\tau_j^{-2} \exp[-(2b^2\tau_j)^{-1}]$ and $(\tau_j + b^{-2})^{-1}$ for Cauchy, Laplace and horseshoe prior respectively. Thus we can see that b controls how the local temperature parameter τ_j distributes. Large b limits the variation

of τ_j and increases the mass around zero, and small b means the range for τ_j to vary is large. This is also consistent to Figure 3 (b), where we can see when $b \rightarrow 0$, τ_j can vary widely, thus the system is in high temperature state, and if $b \rightarrow \infty$, τ_j will be limited around 0 and the system is in low temperature state.

The generalization of Cauchy prior and Laplace prior into the Lévy process mixture not only shows that the connection between Bayesian variable selection and the Ising model with tempering algorithm, it also provides flexibility to choose priors with different shrinkage characteristics. We did not discuss generalization of horseshoe prior as a Lévy process, further discussion can be found in Polson and Scott (2012), but similar conclusion can be drawn for horseshoe prior in terms of shrinkage or tempering.

6 Incorporating Graph Prior Information

In this paper, we mainly discuss model (1) as a graphical model with noninformative prior for γ , and it works well for n is large enough. However, the priori information about γ becomes important when n goes small. There are two purposes of incorporating graph prior information for γ . First, it helps to improve the mixing issue so the model works for $n \ll p$. Second, it improves the power of detecting the true signals. Since two connected nodes with positive interaction intend to be selected or excluded together, only the prior graph for γ with positive interaction is meaningful. If we have the information that some selected nodes and their neighbors are all true nodes, then incorporating a graph prior with those nodes connected will improve the power to identify the nodes with small signal. This is because the prior tells us that those nodes with small signal have more chances to be selected together with their neighbors which are true signal. On the other hand, for those nodes that are not true signal, we have more chances to exclude their neighbors too since the prior tells us they should be excluded together. At first glance, assigning

a prior graph for γ seems like manipulating the weight to select which nodes and their neighbors, but if the prior information is true, then assigning such a prior is reasonable. Even though the prior information is not exactly correct, it will help if the prior graph contains the true graph about which nodes are networked. For example, given a true model, $\mathbf{y} = \sum_{j \in S} \mathbf{x}_j \beta_j$ where $S = 1, \dots, k$ is sequential index up to k and $k < p$. Obviously there are some information about the true variables such that there are k sequential nodes are true nodes, and $p - k$ sequential nodes are not in the true model. Therefore, a Ising prior with one dimensional linear chain will be a very efficient prior since this prior reflects the information that sequential nodes are selected or excluded together. Another example is the genetic pathway data within which different sets of genes function together. Some gene sets are related to the phenotype diseases, some are not. Therefore the prior with this pathway graph helps distinguishing different set of genes in the pathway since among those genes if one node is selected then its connected neighbors have high chance to be selected too. Further example about incorporating prior graph information can be seen in Li and Zhang (2010); Monni and Li (2010); Stingo et al. (2011); Tai et al. (2010).

Since we are only interested in the network prior information, we only apply a graph prior for γ with the interaction matrix $W = \{W_{ij}\}$ without the external field:

$$p(\gamma) \propto \exp \left(\sum_{i < j} W_{ij} \delta_{ij} \right),$$

where W_{ij} represents the prior coupling information between node i and j . For simplicity, considering $W = w\Lambda$, where w is a small positive interaction parameter, and $\Lambda = \{\lambda_{ij}\}$ is the adjacency matrix with $\lambda_{ij} = 1$ if node i and j are connected and $\lambda_{ij} = 0$ if i and j are independent. With this prior, the posterior distribution for γ is modified as:

$$\begin{aligned} p(\gamma|\mathbf{y}, \beta, \phi) &\propto p(\mathbf{y}|\gamma, \beta, \phi) p(\gamma) \\ &\propto \exp \left(\sum_{i < j} J_{ij}^* \delta_{ij} + \sum_j h_j^* \gamma_j \right), \end{aligned} \tag{30}$$

where $J_{ij}^* = (J_{ij} + W_{ij})$. J_{ij} and h_j^* are defined in (10) and (11).

Correspondingly, the two expressions for the cluster algorithm are modified as

$$p_{a,j} = \max \left\{ 1 - \exp \left[(-1)^{\gamma_j} \left(\sum_{k \in c_1} \lambda_{ik} J_{jk}^* - \sum_{l \in c_0} \lambda_{jl} J_{jl}^* \right) \right], 0 \right\}, \quad (31)$$

$$\begin{aligned} & \alpha(\gamma_c^0 \rightarrow \gamma_c^*) \\ &= \min \left\{ \exp \left[\sum_{j \in \bar{c}} (-1)^{\gamma_j} \left(\sum_{k \in c_1} (1 - \lambda_{jk}) J_{jk}^* - \sum_{l \in c_0} (1 - \lambda_{jl}) J_{jl}^* \right) + \sum_{j \in c_0} h_j^* - \sum_{j \in c_1} h_j^* \right], 1 \right\}. \end{aligned} \quad (32)$$

Above two expressions tell us that $p_{a,j}$ and $\alpha(\gamma_c^0 \rightarrow \gamma_c^*)$ are also conditional on Λ .

7 Extension to Nonparametric Regression Models: Bayesian Sparse Additive Model (BSAM)

Although the BVGM is based on the parametric linear regression model (1), it is easy to be extended to nonparametric regression models. Some similar approaches have been suggested, such as nonparametric regression using Bayesian variable selection (Smith and Kohn, 1996) and Bayesian Smoothing Spline ANOVA models (Reich et al., 2009), both use the spline techniques. In the former case, the binary random variable is applied to each knots of spline function in stead of each predictor, thus the model is capable to select the knots of each nonparametric function. In the second paper, second order interactions are included by using function ANOVA. In this paper, we only employ BSAM to demonstrate how easy it is to extend the parametric regression model based on BVGM.

Extending the multiple parametric linear regression model (1) to an additive model is straightforward. In Bayesian point of view, there is no strict difference between parametric and nonparametric additive regression model in sense of that both assign prior to the basis

coefficients. In general, both choosing a basis to express the marginal regression predictor $f_j(\mathbf{x}_j)$. For linear parametric regression, $f_j(\mathbf{x}_j) = \beta_j \mathbf{x}_j$, where the predictor \mathbf{x}_j itself can be considered as the basis to represent f_j and β_j is a univariate random variable. This is just a special case of nonparametric regression model considering $f_j(\mathbf{x}_j) = Z_j \boldsymbol{\beta}_j$, where Z_j is some basis matrix for j th predictor and $\boldsymbol{\beta}_j$ is multivariate random variable, where the basis length $M_j \geq 1$ can vary for different predictor. Despite the variation of the basis chosen, each predictor is corresponding to a univariate random vector $\mathbf{r}_j = f(\mathbf{x}_j) = Z_j \boldsymbol{\beta}_j$. Then the generalized additive model can be expressed as

$$\mathbf{y} = \mu + \sum_{j=1}^p \gamma_j f_j(\mathbf{x}_j) + \epsilon.$$

For this model, similarly, we can consider following prior for $\boldsymbol{\beta}_j$'s

$$\begin{aligned} [\boldsymbol{\beta}_j | \tau_j] &\sim N(\mathbf{0}, b^2 \tau_j^{-1} I), \\ [\tau_j] &\sim p(\tau_j). \end{aligned} \tag{33}$$

where N is multivariate M_j dimensional normal distribution, and $p(\tau_j)$ is some priors similar to previous discussions, such as $G(1/2, 1/2)$, $IG(1, 1/2)$ or $C^+(0, 1)$. For some special basis, the multivariate normal prior of $\boldsymbol{\beta}_j$ may have two variance components such as LS basis (see A.6). Note that because the dimension of $\boldsymbol{\beta}_j$ is changed, if we integrate out τ_j by assigning the same $p(\tau_j)$ as in parametric linear models, the marginal prior $p(\boldsymbol{\beta}_j | b)$ is no longer Cauchy, Laplace, or horseshoe prior any more, but it shares the similar properties as linear parametric case.

Similarly we can define matrix $R = [\mathbf{r}_1, \dots, \mathbf{r}_p]$, design matrices $Z = [Z_1, \dots, Z_p]$, $Z_\gamma = [\gamma_1 Z_1, \dots, \gamma_p Z_p]$ and the coefficients vector $\boldsymbol{\beta} = (\boldsymbol{\beta}_1^T, \dots, \boldsymbol{\beta}_p^T)^T$, but here we should treat $\boldsymbol{\beta}$ and Z as blocks. The total dimension for the design matrix Z is $n \times M$ and $M \times 1$ for $\boldsymbol{\beta}$, where $M = \sum_j M_j$. Without any confusion, we can use the same posterior distribution expressions in (3) and (5) to update $\boldsymbol{\beta}_c$ and ϕ except we need keep in mind $\boldsymbol{\beta}$ and Z are in blocks and

D_c is diagonal block matrix with block $\tau_j/b^2 I_{M_j}, j \in c$ in the diagonal, where I_{M_j} is M_j dimensional identity matrix. Σ_c and $\boldsymbol{\mu}_c$ are expressed as

$$\begin{aligned}\Sigma_c &= (\phi Z_c^T Z_c + D_c)^{-1}, \\ \boldsymbol{\mu}_c &= \phi \Sigma_c Z_c^T (\mathbf{y} - Z_{\gamma_{\bar{c}}} \boldsymbol{\beta}_{\bar{c}}).\end{aligned}\tag{34}$$

The calculation for J_{ij} 's and h_j 's is exactly the same as (10) since those formulas involve R which is a $n \times p$ matrix for both cases. In A.6 we will introduce a specific additive model with the natural cubic spline represented by Lancaster and Šalkauskas (LS) basis. Of course, other spline basis to define Z is possible.

8 Simulation Study

8.1 Case One: Comparison of Three Priors

The first simulation study will examine a simple linear regression model with a general form

$$y = \sum_{j \in S} \beta_j x_j + \epsilon, \tag{35}$$

where $S = \{2, 3, 5, 10\}$, sample size $n = 50$ and $p = 100$, $x_j \sim N(0, 1), j = 1, \dots, p$ and $\epsilon \sim N(0, 1)$. Particularly, we will consider one large signal set and one small signal set: $\{\beta_2, \beta_3, \beta_5, \beta_{10}\} = \{-4, 2, -1, 2.5\}$ and $\{-0.9, 0.7, -0.6, 0.8\}$.

In this simulation, we performed the single site updating with total 6000 iterations for each settings and discarded the first 2000 iterations as burn-in, then calculated the average γ_j 's over total $N = 4000$ iterations as the marginal selection probabilities. Figure 6 plots the marginal selection probability of all variables against the global shrinkage parameter b . For large signals, as shown in the upper row of Figure 6, horseshoe and Cauchy priors perform similarly and show the robustness of large signals, i.e., as b decreases, the selection probability of true signals maintains 1 till very small b and drops to 0.5. Horseshoe prior also shows better robustness than Cauchy prior on the large b side. Both priors have a wide

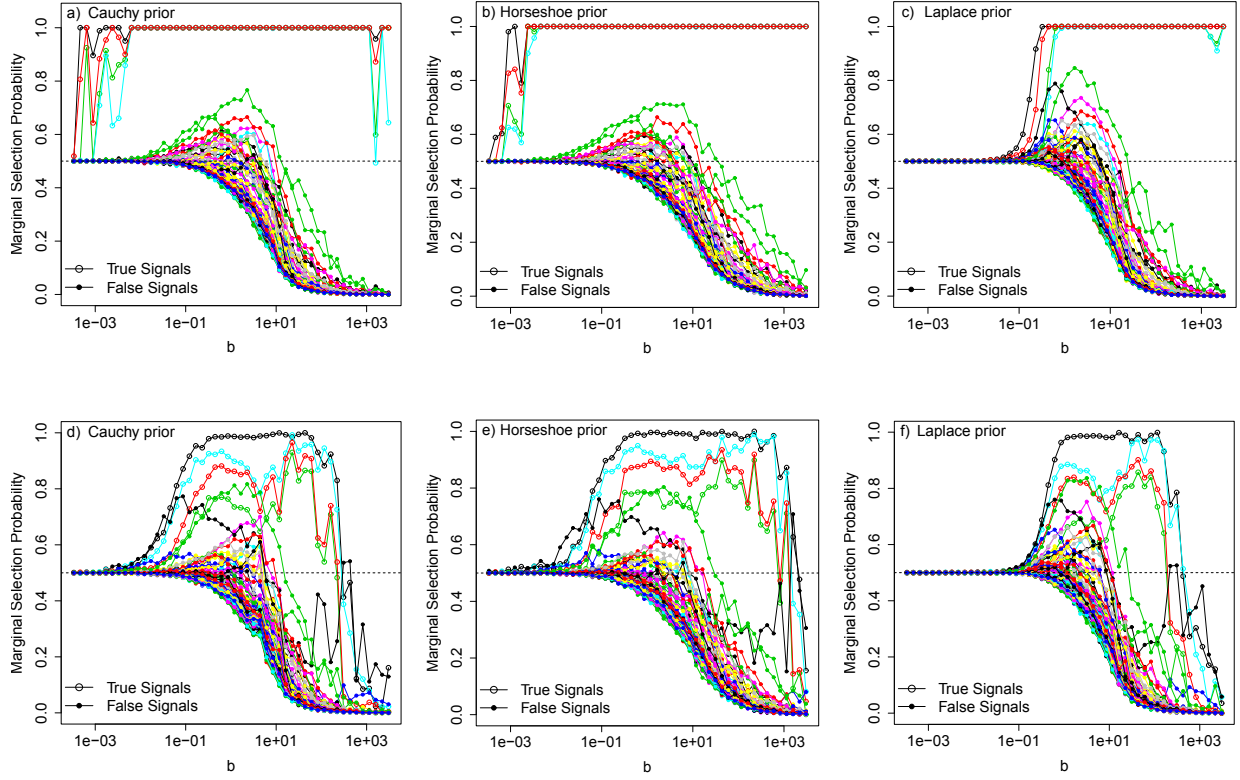


Figure 6: The profile curves of the selection probability of simulation model (35) with different priors for large signal setting (a-c), and small signal setting (d-f).

window of b in which the true signals are well separated from the noise signals. On the other hand, Laplace prior does not demonstrate such robustness for large signals: as the $b \rightarrow 0$, the selection probability of true signals drops to 0.5 very fast. Around $b = 0.1$, all signals reach the 0.5 line for Laplace prior. On the large b side, Laplace prior seems perform a little better than Cauchy prior. Recall the exact calculation of the marginal selection probabilities in Figure 3, we can see that the conclusion made from the simulation about the performance of the three priors is exactly the same.

The bottom row of Figure 6 is the simulation results for small signals. In general, the window for true signals maintaining high selection probability gets narrower for all three

priors. The selection probability of true signal for all priors starts to drop to 0.5 around $b = 0.1$, and around $b = 1000$, they drop to 0. However, the drop rate on both side of b is different for three priors, resulting in different width of the working widow of b . Horseshoe prior has the widest window, Laplace prior gets the narrowest one. Again, the conclusion we made from the simulation is exactly the same as the calculation in Figure 3.

Other observations can be found from Figure 6, especially for small signals. First we can see around $b = 10$, the selection probability of true signals first drops a little and then grows up again. This happens right above the peak of the selection probability of the noise, which indicates potential interaction between the noise and the true signals. This is easy to understand since when sample size is small, the correlation between true signal and noise is large, meaning the parameter ξ_j in (18) is large such that the profile curve is distorted. The second observation from Figure 6 is, although the overall performance of three priors is different in sense of different width of the working window, we can select a right value of b so that for all priors the noise and signals are well distinguishable. For instance, for all priors, with a cut-off probability 0.5 all true signals are separated from the noise at some fixed b between 10 and 1000.

8.2 Case Two: Three Regions of Global Shrinkage Parameter b

Based on the case study one, horseshoe prior has the largest working window, thus in the rest of this paper, we employ horseshoe prior only unless stated otherwise. In this simulation we will examine the case when p is large, say $p = 1000$ or $p = 500$. The linear model still has the form (35) with $x_j \sim N(0, 1), j = 1, \dots, p$ and $\epsilon \sim N(0, 1)$, but we consider following specific models and settings

Model I A. $p = 1000, n = 200, \beta_j = 0.8$ if j is odd; $\beta_j = 1.0$ if j is even. $S = \{31, 91, \dots, 931\} \cup \{60, 120, \dots, 960\}$

B. $p = 1000, n = 500, \beta_j = 0.8$ if j is odd; $\beta_j = 1.0$ if j is even. $S = \{31, 91, \dots, 931\} \cup \{60, 120, \dots, 960\}$

Model II A. $p = 500, n = 100, \beta_j = -0.8$ if j is odd; $\beta_j = 0.8$ if j is even. $S = \{31, 91, \dots, 451\} \cup \{60, 120, \dots, 480\}$

B. $p = 500, n = 500, \beta_j = -0.8$ if j is odd; $\beta_j = 0.8$ if j is even. $S = \{31, 91, \dots, 451\} \cup \{60, 120, \dots, 480\}$

Thus for *Model I* the number of true β_j 's are the cardinality $|S| = 32$, and $|S| = 16$ for *Model II*. For each setting, we performed the single site updating with total 8000 iterations and discarded the first 3000 as burn-in, then calculated the average γ_j 's over total $N = 5000$ iterations as the marginal selection probabilities.

Figure 7 (a-b) and Figure 8 (a-b) plot the marginal selection probability of all variables against b for all settings, so we can have a overall view about all possible global shrinkage. In this simulation it is easier to examine how the working window of b suitable for variable selection changes. For example, in Figure 7 (a), the working window is from $b \approx 0.01$ to $b \approx 2000$, and from $b \approx 0.001$ to $b \approx 3000$ in Figure 7 (b). Within this window, we can see for both setting A and B, b can be further divided into three regions I, II and III. In Figure 7 (a) Region I represents the high temperature or large shrinkage area, where b is around 0.1 or smaller. Region II is a moderate shrinkage area with b between 1 and 100, and the last Region III is around 1000 varying from several hundreds to several thousands. the widths of these three regions also change for different signal strength. We can see in most cases, the signals are well separated from the noise in Region I and III. On the other hand, in Region II, if the signal is not large enough or the sample size n is small, some oscillations or strong interactions occur between signal and noise on the profile curves, resulting in a total mixture up of noise and signal. Thus if to suggest the appropriate value of b , it must be selected to

avoid Region II.

Another interesting observation from this simulation is although in general, both Region I and III both can be used to detect signals, the performance of the MCMC sampling may have different properties in these two regions. We have discussed that Region I has large shrinkage property, but it may not have the sparse consistency. This can be understood from the point of view of the oracle properties of the estimation given the shrinkage parameter. Fan and Li 2001 shows in general the sparse consistency requires $\lambda \rightarrow 0$ or small shrinkage ($\lambda \propto b^{-2}$ is the shrinkage parameter in their paper). Region III, representing small shrinkage area, hence may maintain sparse consistency while Region I loses it. This phenomena is shown in Figure 8 (a) for *Model II A* where the best value of b is in Region III with which all signals are distinguishable from the noise. On the other hand the noise and signals mix up in Region I. This can be seen more clearly in Figure 8 (c-d) for two specific b values: for $b = 223$, most of the true signals have selection probability 1 and distinguished from the noise, while for $b = 0.23$, some the true signals have smaller selection probability than some noise.

How to determine the parameter b is a interesting topic. Some authors suggest assigning another prior for b , such as horseshoe prior Polson and Scott (2011). Unfortunately, because b is a global parameter, when p is large, the posterior distribution of b will be forced to some value that is not in Region III where we prefer. Therefore in this paper, we will not consider assigning a prior for b , instead, we consider it as a tuning parameter. A practical way to select b is to try several b values and choose the one we see the largest gap in selection probability and b is usually between ten to over thousands.




Figure 7: The profile curves of the selection probability of *Model I A* and *B* (a-b). Selection probability at two b values for *Model I A* (c-d).

8.3 Case Three: Comparison of Cluster and Single Site Algorithm

As discussed in Section 5.4, assigning the scale normal mixture prior for β_j 's with shrinkage parameter b is also a tempering algorithm, which means our model already makes improvement in the mixing issue. So there may be no much space left from improving the mixing with a cluster algorithm. We will show that, the performances of cluster and single site algorithms both are b -dependent. In some region of b , one may outperform the other but performs worse in other region.

To demonstrate this, we consider the simple simulation with the same model as (35) with large signals $\{\beta_2, \beta_3, \beta_5, \beta_{10}\} = \{-4, 2, -1, 2.5\}$, $n = 200$ and we vary p from 50 to 1500. We run the simulation with four representative b 's, two are large and two are small, so we can compare the difference behavior of two algorithm with different shrinkage parameters.

To measure the mixing or correlation time, it is convenient to define the “magnetization”, $M^{(i)}$, which represents the average value of the binary random variable γ_j 's at i th sweep of

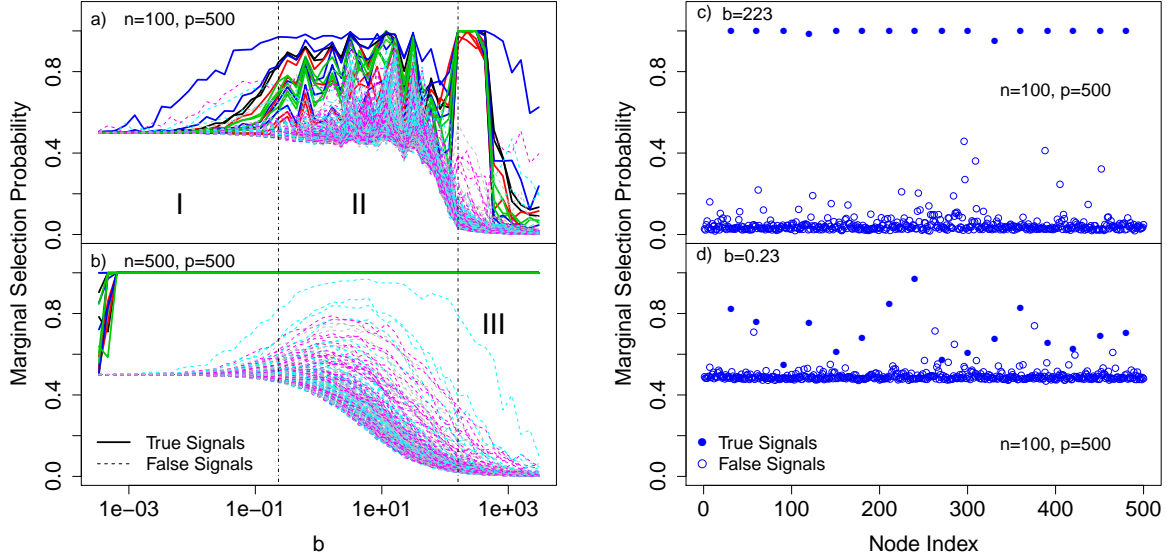


Figure 8: The profile curves of the selection probability of *Model II A* and *B* (a-b). Selection probability at two b values for *Model II A* (c-d).

the MCMC iteration.

$$M^{(i)} = \frac{1}{p} \sum_{j=1}^p \gamma_j^{(i)}.$$

Thus the mixing time of the MCMC iteration can be measured using the time-delayed autocorrelation function (ACF) of the Monte Carlo chain of “magnetization”,

$$C(t) = \frac{\sum_{i=1}^{N-t} (M^{(i)} - \bar{M})(M^{(i+t)} - \bar{M})}{\sum_{i=1}^N (M^{(i)} - \bar{M})^2},$$

where t is the lag or the iteration time from the origin, measured in Monte Carlo sweeps (MCS), and \bar{M} is the average magnetization over total N iterations. We assume the absolute value of $C(t)$ decays exponentially, i.e., $|C(t)| \approx C_0 \exp(-t/\tau)$, where C_0 is some positive constant, and τ is defined as the exponential correlation time. Therefore, we can use τ to measure how fast the chain converges or mixes. The smaller the τ , the faster the system mixes up. Another way to measure the mixing time is simply using the summation of the autocorrelation time, $\sum_{t=0}^L |C(t)|$, where L is the maximum lag calculated.

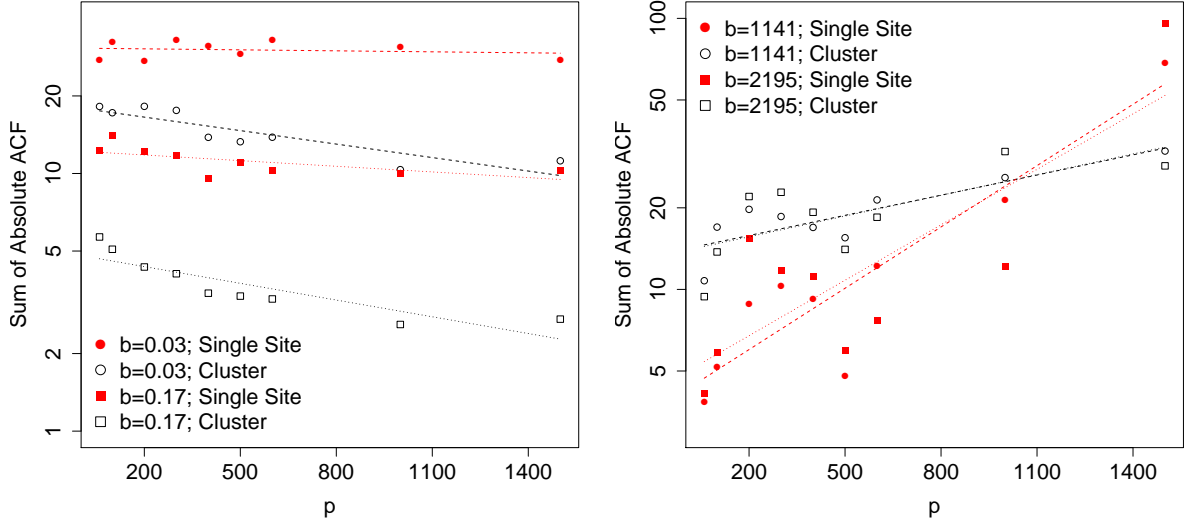


Figure 9: The sum of absolute ACF against variable number p for cluster algorithm and single site algorithm at different b values.

For each p we performed 15000 iterations or sweeps for each setting and discarded the first 5000. From the remaining $N = 10000$ sweeps we calculated the autocorrelation function $C(t)$ up to $L = 100$ lags. Figure 9 shows the summation of absolute ACF time against the nodes size p for $b = 0.03, 0.17, 1141$ and 2195 .

From Figure 9 we can see the different behavior of cluster algorithm and single site algorithm. In large shrinkage region, $b = 0.03$ or 0.17 , the cluster algorithm has mixing time uniformly smaller than the single site algorithm. Note that as the node size increases, the mixing time for all algorithm first decreases slightly and then stabilizes. It may goes up when p goes further. This profile is not well understood yet. Probably because in large shrinkage area, the effect of large node size is pressed by the shrinkage when the number of true nodes is fixed and small. Nevertheless, in this region, we can conclude that cluster algorithm is uniformly outperform the single site algorithm in terms of fast mixing time, and the mixing

time of cluster algorithm is at least two times shorter.

In the small shrinkage region where $b = 1141$ and 2195 , as shown in Figure 9 (b), we see different characteristics. First, the measured mixing time is much more noisy than in Figure 9 (a), but the trend against p is clear. Secondly, unlike large shrinkage area, here we see for both algorithms the mixing time increases as p increase. Furthermore, when p is small, the single site algorithm has shorter mixing time, but slow down very fast as p increases. For example, when $p = 60$, the summation of autocorrelation function is only several MCS, but reaches almost 100 MCS when p is large than 1500, which means extremely slowing down for the MCMC process. On the other hand, although the cluster algorithm is about two times slower when p is small and it also slows down with p increases, the mixing time increases with smaller rate and reaches no more than 50 when $p = 1500$.

Hence in general, we can see cluster algorithm outperforms single site algorithm in terms of mixing time. However, which algorithm should be used depends on the data. Single site algorithm is much less time consuming since the cluster algorithm spends time in forming the cluster. The overall computational time for cluster algorithm is expensive when $p > 1000$. Plus, in many situations, the mixing time may not be so worse for single site algorithm. Thus we prefer using single site algorithm to achieve the results quickly.

8.4 Case Three: Bayesian Sparse Additive Model

In this section, we demonstrate variable selection on following Bayesian sparse additive model:

$$y = f_1(x_1) + f_2(x_2) + f_3(x_3) + f_4(x_4) + \epsilon, \quad (36)$$

where $x_j = (w_j + tu)/(1 + t)$, $j = 1, \dots, p$ and w_1, \dots, w_p and u are iid from Uniform $(0,1)$, and $\epsilon \sim N(0, 1.74)$. Therefore $\text{Corr}(\mathbf{x}_i, \mathbf{x}_j) = t^2/(1 + t^2)$ for $i \neq j$. We consider $t = 0$ and $t = 1$. The later one gives the correlation between two predictors around 0.5. This simulation is

similar to Example 1 in Lin and Zhang (2006) but with $p = 10, 80$ and 150 . We also consider sample size $n = 100$. Functions f_j 's have following forms.

$$\begin{aligned}
f_1(x) &= x, \\
f_2(x) &= (2x - 1)^2, \\
f_3(x) &= \sin(2\pi x)/[2 - \sin(2\pi x)], \\
f_4(x) &= 0.1 \sin(2\pi x) + 0.2 \cos(2\pi x) + 0.3 \sin^2(2\pi x) + 0.4 \cos^3(2\pi x) + 0.5 \sin^3(2\pi x).
\end{aligned} \tag{37}$$

As shown in A.6, for each \mathbf{x}_j , the LS basis employs two precision parameters $\tau_{ej} = \sigma_{ej}^{-2}$ and $\tau_{dj} = \sigma_{dj}^{-2}$. We treat all set of $\{\tau_{ej}, \tau_{dj} : j = 1, \dots, p\}$ independently. Similarly, we can still assign $G(1/2, 1/2)$, $IG(1, 1/2)$, or $C^+(1)$ prior for them. However, since β_j is the $M_j \times 1$ vector and for each node we have two variance components, the marginal prior for β_j given b is no longer simple Cauchy, Laplace or horseshoe prior any more, but it will share the similar properties to its counterpart in linear parametric model. In this simulation, we employ the independent $G(1/2, 1/2)$ prior for each τ_{ej} and τ_{dj} only. For the number of knots of the LS basis, we may consider each predictor has different number of knots, but it turns out a fixed number for all M_j 's, say $M_j = 6$, will give good enough results. Therefore, we fix $M_j = 6, j = 1, \dots, p$ in this simulation. Totally 6000 iterations have been employed by the single site algorithm and first 2000 ones are discarded for all settings.

Table 2: Simulation results of sparse additive model (36) for 500 runs.

	t	p	FP-rate	FN-rate	MS	f_1 SE	f_2 SE	f_3 SE	f_4 SE
BVGM	0	10	0.00(0.02)	0.00(0.03)	3.99(0.17)	0.07(0.05)	0.16(0.06)	0.18(0.08)	0.74(0.27)
		80	0.00(0.01)	0.00(0.03)	4.16(0.48)	0.07(0.05)	0.15(0.06)	0.18(0.08)	0.70(0.27)
		150	0.00(0.02)	0.01(0.04)	4.81(5.52)	0.08(0.09)	0.16(0.07)	0.18(0.11)	0.73(0.44)
	1	10	0.00(0.01)	0.09(0.10)	3.56(0.54)	0.08(0.07)	0.18(0.09)	0.16(0.08)	0.79(0.40)
		80	0.00(0.01)	0.09(0.11)	3.68(0.66)	0.09(0.08)	0.18(0.08)	0.16(0.07)	0.77(0.39)
		150	0.01(0.03)	0.11(0.11)	4.48(7.17)	0.10(0.10)	0.18(0.09)	0.18(0.10)	0.80(0.40)
COSSO	0	10	0.00(0.01)	0.00(0.00)	4.00(0.06)	0.07(0.04)	0.05(0.04)	0.11(0.06)	0.32(0.13)
		80	0.07(0.08)	0.18(0.07)	9.85(11.2)	0.17(0.28)	0.79(0.11)	1.55(0.32)	5.28(0.58)
	1	10	0.01(0.03)	0.04(0.08)	3.86(0.48)	0.07(0.07)	0.26(0.10)	0.14(0.10)	2.00(1.00)
		80	0.10(0.09)	0.19(0.10)	12.5(14.0)	0.36(0.40)	0.37(0.16)	1.05(0.44)	4.67(0.54)




Figure 10: The profile curves of the selection probability of simulation model (36) with $p = 80, n = 100$ for independent setting ($t=0$) (a), and correlated setting ($t=1$) (b).

Figure 10 (a-b) show us how the selection probabilities changes for a range of b with $t = 0$ and $t = 1$ for simulation setting $p = 80$. Note that in Figure 10 (a), one true signal is buried in the noise till $b = 1$, while the same signal is always mixed with noise in Figure 10 (b). As shown in Figure 10, when $b \approx 26$ all false signals go to 0 and we achieve the largest gap between signals and the noise.

One feature of our BVGM is the capability to select the important variables as well as estimate the selected function components at the same time. The four true functions and one noise, the corresponding estimated functions, and the selection probability for all nodes at $b = 26$ of a simulation run with $p = 80$ for $t = 0$ and $t = 1$ are shown in Figure 11 and Figure 12 respectively. Based on the LS basis, the function components estimated are always centered, so we also centered the true functions. In this simulation run, the four true nodes are selected exactly, and the estimated functions of them are calculated by $\hat{f}_j = Z_j E(\beta_j | \gamma_j = 1)$, where the expectation is based on the $N = 4000$ iterations, and




Figure 11: True function f_j 's (blue dashed lines) and estimated function \hat{f}_j 's (blue solid lines) with 95% credible interval (red dashed lines) for the 4 true nodes (a-d) and a noise node (e) of a run of simulation model (36) with independent setting $t = 0$ and $p = 80$. The marginal selection probability at $b = 26$ (f). Note we reordered the first 4 true nodes number to (2, 20, 50, 70) for a better view.

the 95% credible intervals are plotted as well. As shown in Figure 11 and 12, for both $t = 0$ and $t = 1$ the estimated functions are very close to the true functions. Note for the noise function, f_{11} , the selection probability is close to zero, thus the estimated function is calculated by $\hat{f}_j = Z_j E(\beta_j)$, a expectation over all N iterations. This is why we see a very wide credible interval for f_{11} because when $\gamma_j = 0$ the posterior of β_j is multiple normal with large variance. Also note that to have a better view of the selection probability, we reordered the nodes such that the 4 true nodes are (2, 20, 50, 70).

To further examine the performance of variable selection and estimation accuracy of




Figure 12: True function f_j 's (blue dashed lines) and estimated function \hat{f}_j 's (blue solid lines) with 95% credible interval (red dashed lines) for the 4 true nodes (a-d) and a noise node (e) of a run of simulation model (36) with independent setting $t = 1$ and $p = 80$. The marginal selection probability at $b = 26$ (f). Note we reordered the first 4 true nodes number to (2, 20, 50, 70) for a better view.

our method, 500 simulation runs have been employed for $p = 10, 80$ and 100 respectively. We calculated seven statistics: “False Positive Rate (FP-rate)”, “False Negative Rate (FN-rate)”, “Model Size (MS)”, and “Squared Error (SE)” of 4 true functions, where $\text{FP-rate} = \frac{\# \text{False Positive}}{\# \text{False Positive} + \# \text{True Negative}}$, $\text{FN-rate} = \frac{\# \text{False Negative}}{\# \text{False Negative} + \# \text{True Positive}}$, and $\text{SE} = \sum_i^n (f_{j,i} - \hat{f}_{j,i})^2 / n$. The estimated function is calculated by $\hat{f}_j = Z_j E(\beta_j | \gamma_j = 1)$, $j \in \{\text{true nodes}\}$. Since it can happen that $p(\gamma_j = 1 | \mathbf{y}) = 0$ for any true function components, we simply estimate f_j by $\hat{f}_j = 0$ for the 4 true nodes if $p(\gamma_j = 1 | \mathbf{y}) = 0$ in each run. Statistics SE can be used to assess the accuracy of the estimation of the nonlinear function f_j

because the smaller the SE the closer the estimation \hat{f}_j to true function f_j . The average and standard deviation of those statistics over 500 runs are reported in Table 2 and compared with Component Selection and Smoothing Operator (COSSO) (Lin and Zhang, 2006).

As shown in Table 2, the results for our method is pretty robust to p . For each t , all statistics are similar for different p except a little increase in the mean and standardized deviation of those statistics. For different t , our method is also pretty robust, except the increase in the values of FN-rate and SE's. On the other hand, we can see COSSO only performs well for small p . When $p = 80$ (COSSO can not work for $n < p$ case, so no result for $p = 150$), all the statistics of COSSO increase, especially for those four true function components, SE's are very large meaning COSSO can not estimate those function components correctly. In general, we can see our method works very well for BSAM even for large p and large correlation cases in both variable selection and function component estimation.

8.5 Case Four: Linear Chain Prior

Again, we consider the same form of model (35) but with setting $p = 100, n = 100, \beta_j = 0.4$ if j is odd, $\beta_j = 0.8$ if j is even, and $S = \{1, 2, \dots, 15\}$. This example is special in sense of its true predictor set S and false signal set \bar{S} both are continuous in their node index. Obviously, the simplest prior network information is a linear chain: any node's two neighbors are most likely to aligned to this node. Although this is not true for neighbored node 15 and 16, but this discontinuity has small effect on the whole system. With this knowledge, we would consider the linear chain prior for the nodes. $W = \{w\lambda_{ij}\}$, $\lambda_{ij} = 1$ for $|i - j| \leq 1$ and $\lambda_{ij} = 0$ otherwise. In order to have a exchangeable prior, two boundary nodes 1 and p can be treated as neighbors, i.e., $\lambda_{ij} = 1$ if $|i - j| = p - 1$. To fully use this prior information, we employ the cluster algorithm and using this adjacency matrix $\Lambda = \lambda_{ij}$ to form the cluster.

We also compare the results with noninformative prior (employed by single site algorithm).




Figure 13: The graph of a linear chain prior with 20 nodes with 1 through 10 nodes “in” (a). The profile curves of the selection probability of case four model calculated by the cluster algorithm with noninformative prior (b), and with the linear chain prior for γ (c).

Figure 12 (a) shows a example graph of a linear chain with $p = 20$ nodes, note how the two end nodes connected. Figure 13 (b-c) are the selection probability profile plot with noninformative prior and linear chain prior. For each b in both plots, total 6000 iterations have been employed with first 2000 burn-in. For the linear chain prior we take $v = \Phi(\log(b))$ where Φ is the standard normal CDF such that the interaction strength vanishes for large shrinkage and maintains at 1 for small shrinkage. The difference of two plots is obviously: with the noninformative prior, two true signals are very close to the noise for all range of b and hard to be separated from the false signals, while with the linear chain prior, we can see for a large range of b all the true signals are well distinguishable from the noise.

9 Real Data Analysis

9.1 Ozone Data

As an illustration of BSAM implemented by BVGM, we consider an example, the ozone data analyzed by Lin and Zhang (2006). The ozone data is available in R package `cosso`

or **gss**. In the Ozone data, the interesting response variable is the daily maximum one-hour average ozone concentration and eight meteorological variables were recorded in the Los Angeles area for 330 days in 1976. The sample size $n = 330$, and the 8 variables are Height (Hgt), Wind Speed (WS), Humidity (Hum), Temperature (Temp), Inversion Base Height (InvHt), Pressure (Press), Inversion Base Temperature (InvTp), and Visibility (Vis). All predictors were standardized and the response was transformed using logarithm to have normal distributed response. We applied the Bayesian graph model described in Section 5.4 with $M_j = 6$ for all predictors. Total 20000 iterations have been employed with single site algorithm and half of them were discarded as burn in. By quickly examining a series of b , we see the selection probability profile curves are all well defined due to small variable number (not shown). So it is more appropriate to choose a modest shrinkage, which is $b = 1.6$ in this case such that all selected predictors reach their highest selection probability. At $b = 1.6$, two predictors have selection probability less than 0.5.

The estimated results for β_j 's are summarized in Figure 14, where the additive function components, $\hat{f}_j = Z_j E(\beta_j | \gamma_j = 1)$'s, are plotted with 95% credible interval. Because the smallest $P = p(\gamma_j | \mathbf{y})$ is at least 0.14, we have enough iterations for all $\gamma_j = 1$ to estimate all \hat{f}_j . The marginal selection probability, P , for each variable is labeled in each plot. We can then identify three groups of the variables. The first group has at least $P = 0.88$ including Temp, Press, InvHt and Vis. The second group includes InvTp and Hgt with $P = 0.69$ and 0.58. The last group contains Hum and WS with P smaller than 0.5. In variable selection point of view, we will select all variables in the first group surely, and we will not select the last group since their selection probabilities are very close to the baseline. Because of the small variable number and nearly independence of each variable, we can consider the second group as true variables with small signals. In the point view of function components estimation, the selection probability is consistent to its signal estimation. As shown in Figure

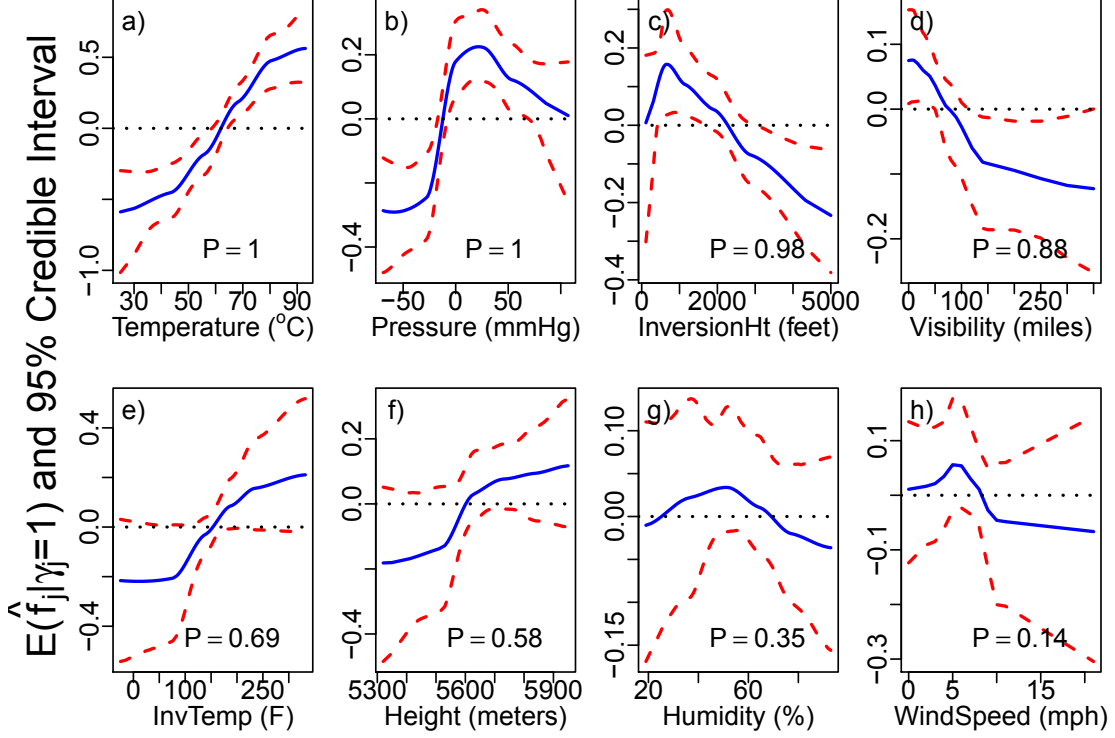


Figure 14: Estimated function \hat{f}_j (blue solid lines) with 95% credible interval (red dashed lines) for the 8 predictors of ozone data labeled by the marginal selection probability $P = p(\gamma_j = 1|\mathbf{y})$ at $b = 1.6$.

16, for the first group of variables the credible intervals only cover a small part of the zero line, while for the third group of variables, the zero line is almost in the center of the credible interval. Although the credible intervals of the second group of variables cover the zero line totally, the zero line is close to the edges of the credible interval.

We also report the summary statistics for all variance components (their inverses) and the intercept in Table 3. In this example we include the intercept term μ in model (36) and assign a prior for μ : $[\mu] \sim N(0, \tau_\mu^{-1})$, and a Gamma prior for τ_μ : $[\tau_\mu] \sim G(4, 2)$. The full conditional distributions for μ and τ_μ are easy to derive (not shown here). Note how the posterior means of these parameters adapt to the data. Especially for τ_{ej} 's and τ_{dj} 's, all start with the same prior, but the posterior means are different. For the first group variables, their

posterior means for τ_{ej} 's and τ_{dj} 's are obviously different from their priors.

Table 3: Parameter estimation of ozone data under Bayesian sparse additive model at $b = 1.6$.

Parameter	Prior		Posterior				
	Mean	Std.Dev.	Mean	Std.Dev.	Median	Lower 2.5%	Upper 2.5%
τ_e^{Temp}	1.000	1.414	0.671	1.192	0.150	0.003	4.199
τ_e^{Press}	1.000	1.414	0.787	1.270	0.282	0.008	4.570
τ_e^{InvHt}	1.000	1.414	0.198	0.597	0.021	0.000	1.700
τ_e^{Vis}	1.000	1.414	1.316	1.517	0.805	0.021	5.485
τ_e^{InvTp}	1.000	1.414	1.010	1.365	0.500	0.010	4.884
τ_e^{Hgt}	1.000	1.414	1.062	1.392	0.547	0.005	5.032
τ_e^{Hum}	1.000	1.414	1.050	1.431	0.509	0.001	5.129
τ_e^{WS}	1.000	1.414	1.104	1.505	0.545	0.001	5.334
τ_d^{Temp}	1.000	1.414	0.226	0.724	0.026	0.002	2.212
τ_d^{Press}	1.000	1.414	0.033	0.047	0.021	0.002	0.125
τ_d^{InvHt}	1.000	1.414	0.889	1.352	0.349	0.004	4.731
τ_d^{Vis}	1.000	1.414	0.751	1.235	0.227	0.004	4.226
τ_d^{InvTp}	1.000	1.414	1.122	1.413	0.611	0.005	5.227
τ_d^{Hgt}	1.000	1.414	1.045	1.429	0.521	0.009	5.140
τ_d^{Hum}	1.000	1.414	1.061	1.474	0.513	0.001	5.048
τ_d^{WS}	1.000	1.414	0.949	1.383	0.404	0.000	4.985
ϕ	-	-	6.602	0.542	6.585	5.599	7.681
μ	0.000	$\sqrt{\tau_\mu}$	2.143	0.066	2.145	2.011	2.265
τ_μ	2.000	1.000	1.043	0.495	0.971	0.316	2.205

9.2 Gene Selection in Pathway Data

Mootha et al. (2003) presented an pathway based analysis to test a priori defined pathways for association with the diabetes disease. A pathway is a predefined set of genes that serve a particular cellular or physiological function. Therefore a genetic pathway can be expressed by a graph to prrsent the gene network within this pathway. Mootha et al. (2003) identified several significant pathways among which ‘‘Oxidative phosphorylation’’, ‘‘Alanine-and-aspartate metabolism’’ et al. are interesting ones. However, even with those significant pathways identified, gene selection in microarray data analysis is still difficult because alterations in gene expression are modest due to the large number of genes, small sample sizes and variability between subjects. Stingo et al. (2011) provide a Bayesian technique to

incorporate biological information into linear models to select genes and pathways. Similar to Stingo et al. (2011), we also incorporate the pathway network information into our graph model, and apply it to gene selection of the diabetes data from Mootha et al. (2003). However, in our method, we use the gene network information in the pathway as the prior for γ , and we don't select pathways. The data contains gene expressions from $n = 35$ subjects, 17 normal and 18 Type II diabetes patients. We merged three interesting pathways, "Oxidative phosphorylation", "Alanine-and-aspartate metabolism" and "Glutamate-metabolism" into one graph with total $p = 173$ nodes (some nodes are different probe sets of the same gene, so the gene names are identical) which is a subgraph of the corresponding merged graph obtained from KEGG database. The response \mathbf{y} is the continuous glucose level.




Figure 15: Profile curves of the selection probability of genetic pathway data with noninformative prior for γ (a), and with informative prior as (38) (b).

The top left plot of Figure 16 shows the network of our merged gene set. Note, the prior required for our graph model is undirected graph with positive interaction only. We can see, most of the nodes are independent in this data set, and there are only three genetic clusters.

Because of this, if we apply the cluster algorithm and use the adjacency matrix Λ based this network information into expression (31), we will end up with a few nodes in the same genetic cluster that can form the clusters for the algorithm. Therefore, we consider following interaction matrix $W = \{w_{ij}\lambda_{ij}\}$ for the prior of γ with adjacency matrix $\Lambda = \{\lambda_{ij}\}$ as

$$\begin{aligned} \lambda_{ij} &= 1, \quad i, j = 1, \dots, p \\ w_{ij} &= \begin{cases} w, & i \notin S \text{ or } j \notin S \\ w + \Delta w, & i \in S \text{ and } j \in S, \end{cases} \end{aligned} \quad (38)$$

where S represents one of the three genetically networked gene clusters in the pathway network, w and Δw are small positive numbers stand for the strength of the interaction in the prior and the difference of two types of interaction. If $\Delta w = 0$, we can consider (38) as a baseline graph prior for γ , which is a complete graph with positive fixed interaction. Since we also vary b to have an overall view about the selection probability, it is necessary to have $w \rightarrow 0$ when $b \rightarrow 0$ since with large shrinkage $J_{ij} \rightarrow 0$ and we don't want w_{ij} dominates the graph interaction. One convenient way is to express w as $w = w_0 \Phi[\log(b)]$ which approaches 0 as $b \rightarrow 0$ and reaches the maximum w_0 for large b , where $\Phi(\cdot)$ is the CDF of standard normal. Note that the choice of w is involved in the consideration of so called phase transition (Li and Zhang, 2010). If w is too large all the nodes will always be connected which leads to either all nodes are selected or none are selected. Now we consider $\Delta w \neq 0$, say $\Delta w = 5w$, so we incorporate the genetic network information into the graph prior. Δw can not be too large, otherwise those genes in the genetic clusters will always be aligned which means in this data set they will all have small selection probability. So we choose the $w_0 \approx 0.01$ as small as possible to avoid the phase transition phenomena, but it is must be large enough to reduce so called region II of b caused by small sample size. With this selection, we have the prior interaction $w_{ij} \approx 0.01$ for two nodes not in the genetic cluster together for large b , and $w_{ij} \approx 0.06$ for two nodes in the cluster for large b .

As shown in Figure 15, the effect of incorporating prior information for the graph model is obvious. We run the cluster algorithm for total $N = 40000$ iterations, and discarded the first 10000 as burn-in. So the selection probability is calculated by taking the mean of γ over 30000 iterations. In Figure 15 (a), with noninformative prior for γ , we can see even though we are still able to identify several genes behaving differently from the rest (highlighted by solid lines in the plot), for the moderate value of b all the curves are mixed. On the other hand, in Figure 15 (b), with informative prior for γ defined as (38) the profile curves are much “cleaner” even for moderate b . Around $b = 10$ we can see a bunch of curves are clearly distinguishable from the rest. We highlighted 6 nodes with highest selection probability around $b = 10$ in Figure 15 (b).

To examine more details of the results, in Figure 16 we fixed $b = 8.5$ and run the cluster algorithm for $N = 60000$ iterations with first 20000 discarded. With this shrinkage parameter, the prior interaction parameter $w_{ij} \approx 0.06$ for i and j in the genetic cluster, and $w_{ij} \approx 0.01$ otherwise. The selection probability for all nodes are shown in bottom left of Figure 16 where we take a cut-off probability as 0.2 and identify 6 nodes that have relative high selection probabilities. Among those nodes, UQCRB has the largest selection probability for all range of b , so it is easy to identify UQCRB as the most significant gene. We also select other five genes, COX8, ATP5G2 (two probe sets), ATP5H and CRAT at $b = 8.5$. All the genes selected except CRAT are from “Oxidative phosphorylation” pathway which is related to ATP synthesis. It is well known ATP plays a importance role in Type II diabetes disease. CRAT is from “Alanine-and-aspartate metabolism” pathway. Both “Oxidative phosphorylation” and “Alanine-and-aspartate metabolism” pathway are two top significant pathways identified using random forrest tree approach (Pang et al., 2006).

Since our cluster algorithm forms the Wolff cluster at each iteration, a byproduct of the MCMC sampler is the frequency of two nodes being aligned or anti-aligned when they are

in the cluster. The top right plot in Figure 16 is the heatmap matrix of the frequency of two nodes being aligned in the cluster out of 40000 iteration, and bottom left plot is the frequency of two nodes being anti-aligned in the cluster. The color bar of two plots shows the scale of the frequency, the darker the color the lower the frequency. In the top right plot, the dark colored lines are those genes have lower chance to be aligned to the others when they form the cluster, and in the bottom left plot, the bright colored lines are the same genes but with high chance to be anti-aligned to others if they form the cluster. Note those lines are consistent to the genes with high selection probability in the bottom right plot. This is because most of the genes have low selection probability around 0.05 then those genes with higher selection probabilities should have lower (higher) chance to be (anti-)aligned with them. For individual node, we define it is always self-aligned, so the diagonal in top right plot is 1, meanwhile an individual node is never anti-aligned to itself, so diagonal in bottom left plot has value 0. The distinguishable color of those genes in two heatmaps show that we can also use the cluster information to distinguish genes.

So far, we identify 6 genes (probe sets) with cut-off probability 0.2, we may decrease the cut-off to select more genes. However, the selection probabilities are low for most of the genes except for UQCRB at fixed b . This is because of the problem of modest alterations for single gene selection, or it simply means the signals are weak. Here we selected those genes not only depending on the selection probability at fixed b , in stead we select them by examining their overall profile as shown in Figure 15. We also demonstrated that the graph model variable selection can easily adopt the prior graph information, thus we can consider similar approach as Stingo et al. (2011) to select networked pathways, which may result in higher selection probability for pathways at the optimal shrinkage parameter b .




Figure 16: Summary of the results for the genetic pathway data. Top left: genetic network structure of the data. Top right: the frequency matrix of two nodes aligned in the cluster over total iterations. Bottom left: the frequency matrix of two nodes anti-aligned in the cluster over total iterations. Bottom right: Selection probability with cluster algorithm at $b = 8.5$ with informative prior (38).

10 Discussion

The goal of this paper is to present BVGM from two major aspects. The first is how to sample the “in” or “out” binary random variable. We pointed out that Bayesian variable selection can be considered as the binary random process on a complete graph given noninformative prior for γ_j 's, and we compared the single site and generalized Wolff cluster updating algorithm. Another one is how to construct the interaction matrix of the complete graph, which is implemented by sampling the linear model coefficient β_j 's through the scale mixtures of normal priors. We also discussed the marginal selection probability profile under different shrinkage parameter and compared three prior settings for β_j which represent three typical situations of shrinkage proportion. Our BVGM method possesses the advantages of simple form, easy implementation and straightforward to extension. For example, the BVGM is very easy to extend to Bayesian sparse model by representing the nonparametric function components f_j as linear combination of the basis matrix $f_j = Z_j\beta_j$, then we can employ the group selection of vector β_j . Another example is to incorporate network information for γ . Although this paper does not focus on how to construct the prior network structure information, the simulation and real data analysis show that it is easy to incorporating the prior graph information and improve the performance of BVGM. This paper also systematically studies the behaviors of the marginal selection probability against the shrinkage. Both theoretical and simulated results show that to have the largest gap between the signals and the noise it is critical for the scale mixture normal prior to maintain substantial proportion on small shrinkage.

However, this paper only starts a different view angle about Bayesian variable selection, further research includes but are not limited to following questions.

1. As shown in Theorem 3 we have $E(\beta_j|\mathbf{y}, \gamma_j = 1) = \frac{d}{da_j} \log \pi_j^b$ for orthogonal design.

This equation reveals the relationship between the selection probability and the signal magnitude: the larger the signal, the further the selection probability profile being separated from the baseline. However, this relationship does not provides a cut-off rule to separate the signals from the noise. As we can see in the paper, at different shrinkage parameter b , the selection probabilities are different. Thus at different shrinkage parameter, the cut-off line should be different too, and we simply choose b where there is a largest gap between two set of signals. However, $p(\gamma_j|\mathbf{y}) > 0.5$ does not mean the corresponding predictor should be selected, such as for $b \rightarrow 0$ many noise predictors have selection probability no less than 0.5; and $p(\gamma_j|\mathbf{y}) \rightarrow 0$ does not guarantee the corresponding predictor should be removed, since all predictors have zero selection probability for very large b with orthogonal design. Further research should show the consistency of selecting the predictors based on the “largest gap” rule, and provides more straightforward method to choose b .

2. Limitation of fixing b . The global shrinkage or temperature parameter b is fixed, which limits the performance of our method since it limits the range of local shrinkage parameter. Assigning prior for b is not appropriate too due to the high dimensionality such that the posterior distribution of b is forced to be very small. To automatically have b large and small at the same time, we can adopt a remedy similar to exchange Monte Carlo by running parallel MCMC’s at two or more b ’s with some b small and the other large, and exchanging their configuration according to certain probability satisfying the detailed balance. This remedy may improve the performance of results.
3. Bayesian sparse additive model with interaction. Our BSAM does not include the interaction terms, but it should be easy to extend to include them. The only problem is figure out how to represent the interaction function components. This can be done

under the spline ANOVA models similar to Reich et al. (2009).

4. Prior graph information. With a known networked graph prior, we have better performance in some cases because the prior reduces the searching space for γ . However, there is no way to have exact knowledge about the network prior information for the predictors, and it is difficult to construct a meaningful network as the prior. So this keeps a open question as discussed by Li and Zhang (2010) and Monni and Li (2010).

Appendix A

A.1 Proof of Theorem 1

In general, in the Markov chain of MH algorithm, the move from current state γ_c^0 to the proposed state γ_c^* in the cluster has the transition probability, $P(\gamma_c^0 \rightarrow \gamma_c^*)$, which satisfies the detailed balance condition

$$\frac{P(\gamma_c^0 \rightarrow \gamma_c^*)}{P(\gamma_c^* \rightarrow \gamma_c^0)} = \frac{p(\gamma_c^* | \mathbf{y}, \beta, \phi)}{p(\gamma_c^0 | \mathbf{y}, \beta, \phi)} = \exp \{ - [U(\gamma_c^*) - U(\gamma_c^0)] \}. \quad (\text{A.1.1})$$

The transition probability can be broken down into two parts:

$$P(\gamma_c^0 \rightarrow \gamma_c^*) = g(\gamma_c^0 \rightarrow \gamma_c^*) A(\gamma_c^0 \rightarrow \gamma_c^*),$$

where $g(\cdot)$ is the selection probability, which is the probability given γ_c^0 that the new target state generated, and $A(\cdot)$ is the acceptance ratio. Thus

$$\frac{g(\gamma_c^0 \rightarrow \gamma_c^*) A(\gamma_c^0 \rightarrow \gamma_c^*)}{g(\gamma_c^* \rightarrow \gamma_c^0) A(\gamma_c^* \rightarrow \gamma_c^0)} = \exp \{ - [U(\gamma_c^*) - U(\gamma_c^0)] \}. \quad (\text{A.1.2})$$

Now we consider the move $\gamma_c^0 \rightarrow \gamma_c^*$, starting with a particular cluster c and then adding the others to it in a particular order. Consider also the reverse move, which takes us back to γ_c^0 from γ_c^* , starting with exactly the same cluster (except the state in the cluster is flipped), and

adding the others to it in exactly the same way as in the forward move. The probability of choosing the cluster (if the cluster is the seed node) is exactly the same in the two directions, as is the probability of adding each node to the cluster. The only difference between the two directions is the probability of “breaking” bonds around the edge of the cluster. Because the cluster couples with all $j \in \bar{c}$, for both directions, there are $|\bar{c}|$ bonds which have to be broken in order to flip the cluster. These broken bonds represent the affinity between the cluster and the spins which were not added to the cluster by the algorithm. We represent the probability of not adding such a node in forward move as $1 - p_{a,j}^0, j \in \bar{c}$ and in backward move as $1 - p_{a,j}^*, j \in \bar{c}$. Thus the probability of not adding all of them, which is proportional to the selection probability $g(\gamma_c^0 \rightarrow \gamma_c^*)$ for the forward move, is $\prod_{j \in \bar{c}} (1 - p_{a,j}^0)$. In the reverse move then the probability of doing it is $\prod_{j \in \bar{c}} (1 - p_{a,j}^*)$. The condition of detailed balance, Equation (A.1.1), along with Equation (A.1.2), then tells us that

$$\frac{g(\gamma_c^0 \rightarrow \gamma_c^*) A(\gamma_c^0 \rightarrow \gamma_c^*)}{g(\gamma_c^* \rightarrow \gamma_c^0) A(\gamma_c^* \rightarrow \gamma_c^0)} = \prod_{j \in \bar{c}} \left(\frac{1 - p_{a,j}^0}{1 - p_{a,j}^*} \right) \cdot \frac{A(\gamma_c^0 \rightarrow \gamma_c^*)}{A(\gamma_c^* \rightarrow \gamma_c^0)} = \exp \left\{ - [U(\gamma_c^*) - U(\gamma_c^0)] \right\}. \quad (\text{A.1.3})$$

Note that the energy change $U(\gamma_c^*) - U(\gamma_c^0)$ is only determined by the bonds (the coupling between c and \bar{c}) and coupling of c with the external field \mathbf{h}^* , i.e.,

$$U(\gamma_c^*) - U(\gamma_c^0) = \sum_{j \in \bar{c}} (-1)^{\gamma_j} \left(\sum_{k \in c_0} J_{jk} - \sum_{l \in c_1} J_{jl} \right) + \sum_{j \in c_1} h_j^* - \sum_{j \in c_0} h_j^*. \quad (\text{A.1.4})$$

The first part of right hand side of Equation (A.1.4) can be decomposed as

$$\lambda \sum_{j \in \bar{c}} (-1)^{\gamma_j} \left(\sum_{k \in c_0} J_{jk} - \sum_{l \in c_1} J_{jl} \right) + (1 - \lambda) \sum_{j \in \bar{c}} (-1)^{\gamma_j} \left(\sum_{k \in c_0} J_{jk} - \sum_{l \in c_1} J_{jl} \right).$$

With the probability of adding a node $j \in \bar{c}$ to the cluster, $p_{a,j}$, defined as (15),

$$\frac{1 - p_{a,j}^0}{1 - p_{a,j}^*} = \exp \left\{ -\lambda (-1)^{\gamma_j} \left(\sum_{k \in c_0} J_{jk} - \sum_{l \in c_1} J_{jl} \right) \right\}$$

Substituting above equation into Expression (A.1.3) and rearranging, we derive the acceptance ratio for the moves in the two directions as

$$\frac{A(\gamma_c^0 \rightarrow \gamma_c^*)}{A(\gamma_c^* \rightarrow \gamma_c^0)} = \exp \left[(1 - \lambda) \sum_{j \in \bar{c}} (-1)^{\gamma_j} \left(\sum_{k \in c_1} J_{jk} - \sum_{l \in c_0} J_{jl} \right) + \sum_{j \in c_0} h_j^* - \sum_{j \in c_1} h_j^* \right],$$

and the acceptance probability for move from γ_c^0 to γ_c^* is

$$\alpha(\gamma_c^0 \rightarrow \gamma_c^*) = \min \left\{ \frac{A(\gamma_c^0 \rightarrow \gamma_c^*)}{A(\gamma_c^* \rightarrow \gamma_c^0)}, 1 \right\}.$$

As well as satisfying the detailed balance, the algorithm also guarantees the ergodicity by the fact that there is always a finite chance that any spin will be chosen as the sole member of cluster of one, which is then flipped. The appropriate succession of such moves will get us from any state to any other in a finite time as ergodicity requires.

A.2 Proof of Theorem 2

The proof of the first part in Theorem 2 is simply algebra calculation. First define $\mathbf{y}^* = \mathbf{y} - \sum_{k \neq j} \gamma_k \mathbf{x}_k \beta_k$, and integrate out β_j and $\beta_{\bar{j}}$ separately in following expression

$$\begin{aligned} p(\gamma|\mathbf{y}, \boldsymbol{\tau}, b) &\propto \int p(\mathbf{y}|\gamma, \boldsymbol{\beta}) p(\boldsymbol{\beta}|\boldsymbol{\tau}, b) d\boldsymbol{\beta} \\ &\propto \int \exp \left[-\frac{1}{2} (\mathbf{y}^* - \gamma_j \mathbf{x}_j \beta_j)^2 \right] p(\beta_j|\tau_j, b) d\beta_j \prod_{k \neq j} p(\beta_k|\tau_k, b) d\beta_k \\ &\propto \exp \left[\frac{\gamma_j}{2} a_j^2 \left(\gamma_j + \frac{\tau_j}{b^2} \right)^{-1} \right] \left(\frac{\tau_j/b^2}{\gamma_j + \tau_j/b^2} \right)^{1/2} \cdot \xi(\gamma_j, \kappa_j, \gamma_{\bar{j}}, \boldsymbol{\tau}_{\bar{j}}), \end{aligned} \quad (\text{A.2.1})$$

where $\xi(\gamma_j, \kappa_j, \gamma_{\bar{j}}, \tau_{\bar{j}})$ is calculated by integrating out $\beta_{\bar{j}}$:

$$\begin{aligned}
\xi(\gamma_j, \kappa_j, \gamma_{\bar{j}}, \tau_{\bar{j}}) &\propto \int \exp \left[\frac{\gamma_j}{2} \left(\gamma_j + \frac{\tau_j}{b^2} \right)^{-1} \left(\sum_{k \neq j; l \neq j} \beta_k \gamma_k \mathbf{x}_k^T \mathbf{x}_j \mathbf{x}_j^T \mathbf{x}_l \beta_l \gamma_l - 2 \sum_{k \neq j} \mathbf{x}_k^T \mathbf{x}_j \mathbf{x}_j^T \mathbf{y} \beta_k \gamma_k \right) \right] \\
&\quad \times \exp \left(-\frac{\mathbf{y}^{*T} \mathbf{y}^*}{2} \right) \prod_{k \neq j} p(\beta_k | \tau_k, b) d\beta_k \\
&\propto \int \exp \left[\frac{\gamma_j}{2} \left(\gamma_j + \frac{\tau_j}{b^2} \right)^{-1} \left(\beta_{\bar{j}}^T (\mathbf{c}_{\gamma_{\bar{j}}} \mathbf{c}_{\gamma_{\bar{j}}}^T) \beta_{\bar{j}} - 2a_j \mathbf{c}_{\gamma_{\bar{j}}}^T \beta_{\bar{j}} \right) \right] \\
&\quad \times \exp \left[-\frac{1}{2} \left(\mathbf{y} - X_{\gamma_{\bar{j}}} \beta_{\bar{j}} \right)^2 \right] p(\beta_{\bar{j}} | \tau_{\bar{j}}, b) d\beta_{\bar{j}} \\
&\propto \exp \left[\frac{1}{2} \left(\mathbf{a}_{\bar{j}} - (1 - \kappa_j^{\gamma_j}) a_j \mathbf{c}_{\gamma_{\bar{j}}} \right)^T \Omega_j^{-1} \left(\mathbf{a}_{\bar{j}} - (1 - \kappa_j^{\gamma_j}) a_j \mathbf{c}_{\gamma_{\bar{j}}} \right) \right] |\Omega_j|^{1/2} |D_{\bar{j}}|^{1/2},
\end{aligned} \tag{A.2.2}$$

where $\Omega_j = [D_{\bar{j}} + X_{\gamma_{\bar{j}}}^T X_{\gamma_{\bar{j}}} - (1 - \kappa_j)^{\gamma_j} (\mathbf{c}_{\gamma_{\bar{j}}} \mathbf{c}_{\gamma_{\bar{j}}}^T)]^{-1}$ and is easy to show it is positive definite.

We also used the identity $\gamma_j / (\gamma_j + \frac{\tau_j}{b^2}) = (1 - \kappa_j^{\gamma_j})$. Note if $\gamma_j = 0$, $\xi(\dots)$ does not depend on κ_j or τ_j .

Then by definition and above expressions,

$$\begin{aligned}
\pi_j^b &= \frac{\int \sum_{\gamma_{\bar{j}}} p(\gamma_j = 1, \gamma_{\bar{j}} | \mathbf{y}, \boldsymbol{\tau}, b) p(\boldsymbol{\tau}) d\boldsymbol{\tau}}{\int \sum_{\gamma_{\bar{j}}} p(\gamma_j = 0, \gamma_{\bar{j}} | \mathbf{y}, \boldsymbol{\tau}, b) p(\boldsymbol{\tau}) d\boldsymbol{\tau}} \\
&= \int \pi_j \cdot \frac{\int \sum_{\gamma_{\bar{j}}} \xi(\gamma_j = 1, \kappa_j, \gamma_{\bar{j}}, \tau_{\bar{j}}) p(\tau_{\bar{j}}) d\tau_{\bar{j}}}{\int \sum_{\gamma_{\bar{j}}} \xi(\gamma_j = 0, \kappa_j, \gamma_{\bar{j}}, \tau_{\bar{j}}) p(\tau_{\bar{j}}) d\tau_{\bar{j}}} \cdot p(\tau_j) d\tau_j \\
&= \int \pi_j \xi_j p(\kappa_j) d\kappa_j
\end{aligned} \tag{A.2.3}$$

with x_j defined as (20). It is easy to show that $\xi_j = 1$ for orthogonal design since $\mathbf{c}_{\gamma_{\bar{j}}} = \mathbf{0}$ for all j .

The proof of the second part for π_j is trivial. Obviously, $\pi_j \rightarrow 1$ as $\kappa_j \rightarrow 1$ and $\pi_j \rightarrow 0$ as $\kappa_j \rightarrow 0$. For π_j^b , it is more convenient using $\pi_j^b = \int \pi_j \xi_j p(\tau_j) d\tau_j$, where π_j and ξ_j are measurable functions indexed by b . Both π_j and ξ_j are bounded by some positive number for all b . When $b \rightarrow 0$, $\lim_{b \rightarrow 0} \pi_j \rightarrow 1$ and $\lim_{b \rightarrow 0} \xi_j \rightarrow 1$, thus according to Lebesgue's Dominated Convergence Theorem (DCT), $\lim_{b \rightarrow 0} \pi_j^b = \lim_{b \rightarrow 0} \int \pi_j \xi_j p(\tau_j) d\tau_j = \int \lim_{b \rightarrow 0} (\pi_j \xi_j) p(\tau_j) d\tau_j = 1$. When $b \rightarrow \infty$, $\lim_{b \rightarrow \infty} \pi_j \rightarrow 0$ and $\lim_{b \rightarrow \infty} \xi_j$ equal to some finite number. Again the limit

and integral commute by DCT, thus we have $\lim_{b \rightarrow \infty} \pi_j^b = \int \lim_{b \rightarrow \infty} (\pi_j \xi_j) p(\tau_j) d\tau_j = 0$.

A.3 Proof of Theorem 3

The existence of m_j indicates the marginal prior $p(\beta_j) = \int p(\beta_j | \tau_j) p(\tau_j) d\tau_j$ is bounded for $\beta_j \in \mathbb{R}$, which is true for Cauchy and Laplace prior. Using identity

$$(\beta_j - a_j) p(\mathbf{y} | \beta_j, \gamma_j = 1) = \mathbf{x}_j^T \frac{\partial}{\partial \mathbf{y}} p(\mathbf{y} | \beta_j, \gamma_j = 1),$$

so that

$$\begin{aligned} m_j [E(\beta_j | \mathbf{y}, \gamma_j = 1) - a_j] &= \int (\beta_j - a_j) p(\mathbf{y} | \beta_j, \gamma_j = 1) p(\beta_j) d\beta_j \\ &= \int \mathbf{x}_j^T \frac{\partial}{\partial \mathbf{y}} p(\mathbf{y} | \beta_j, \gamma_j = 1) p(\beta_j) d\beta_j \\ &= \mathbf{x}_j^T \frac{\partial}{\partial \mathbf{y}} \log m_j. \end{aligned} \tag{A.3.1}$$

Following the lemma given in Pericchi and Smith (1992), the interchange of the derivative and the integral is justified. The second result of Theorem 3 is straightforward by observing $m_j \propto \exp(-\frac{1}{2} \mathbf{y}^T \mathbf{y}) \pi_j^b$, thus

$$\mathbf{x}_j^T \frac{\partial}{\partial \mathbf{y}} \log m_j = \mathbf{x}_j^T \left(-\mathbf{y} + \frac{d \log \pi_j^b}{da_j} \frac{da_j}{d\mathbf{y}} \right) = -a_j + \frac{d \log \pi_j^b}{da_j}.$$

For horseshoe prior, $p(\beta_j)$ is not bounded. However, using the technique introduced in Carvalho and Polson (2010) by defining $m_j^* = \int p(\mathbf{y} | \beta_j, \gamma_j = 1) p(\beta_j | \tau_j) p(\tau_j) \tau_j^{-1} d\beta_j d\tau_j$, it can be shown

$$E(\beta_j | \mathbf{y}, \gamma_j = 1) = -\frac{m_j^*}{m_j} \mathbf{x}_j^T \frac{\partial}{\partial \mathbf{y}} \log m_j^*,$$

and similar arguments then follow. For horseshoe prior, it also can be shown

$$\frac{\pi_j^b}{\pi_j^{b*}} \frac{d}{da_j} \log \pi_j^b = \frac{1}{\pi_j^{b*}} \frac{d}{da_j} \left(\int \pi_j (\tau_j^{-1} + 1) p(\tau_j) d\tau_j \right) - E(\kappa_j | \mathbf{y}, \gamma_j = 1) a_j$$

where $\pi_j^{b*} = \int \pi_j \tau_j^{-1} p(\tau_j) d\tau_j$.

A.4 Proof of Theorem 4

It is more convenient to use following equivalent representation of π_j^b

$$\pi_j^b = \int \exp \left[\frac{a_j^2}{2} (1 + \tau_j)^{-1} \right] [\tau_j / (1 + \tau_j)]^{\frac{1}{2}} p(\tau_j | b) d\tau_j,$$

where $p(\tau_j | b)$ is corresponding prior of τ_j given b such that $\pi_j^b = \int p(\beta_j | b) p(\tau_j) d\tau_j = \int p(\beta_j) p(\tau_j | b) d\tau_j$. Then the similar condition for $p(\tau_j | b)$ can be derived from the condition for $p(\tau_j)$ in Theorem 4, i.e.,

$$p(\sigma_j^2 | b) \sim (\sigma_j^2)^{\alpha-1} \exp \left(-\frac{\lambda}{b^2} \sigma_j^2 \right) L^b(\sigma_j^2) d\sigma_j^2, \text{ as } \sigma^2 \rightarrow \infty, \quad (\text{A.4.1})$$

where L^b is the slowly varying function conditioning on parameter b .

Then the marginal odds π_j^b can be expressed as

$$\pi_j^b \propto \exp \left(\frac{a_j^2}{2} \right) m^b(a_j) = \exp \left(\frac{a_j^2}{2} \right) \int \omega_j^{-1} \exp \left(-\frac{a_j^2}{\omega_j^2} \right) p(\omega_j^2 | b) d\omega_j^2,$$

where $\omega_j^2 = 1 + \sigma_j^2$, and the integral, $m^b(a_j)$, is a scale mixture of normals. Now the proof is similar to Polson and Scott (2011). If prior $p(\sigma_j^2 | b)$ satisfies the conditions defined in (A.4.1), so does $p(\omega_j^2 | b)$ satisfy similar conditions, i.e.,

$$p(\omega_j^2 | b) \sim (\omega_j^2)^{\alpha-1} \exp \left(-\frac{\lambda}{b^2} \omega_j^2 \right) L^b(\omega_j^2) \text{ as } \omega_j^2 \rightarrow \infty$$

Then following Theorem 6.1 of Barndorff-Nielsen et al. (1982), as $a_j \rightarrow \infty$, $m^b(a_j)$ can be approximated as

$$m^b(a_j) \sim \begin{cases} |a_j|^{2\alpha-1} L^b(a_j^2) & \text{if } \lambda = 0 \\ |a_j|^{\alpha-1} \exp \left(-\sqrt{\frac{2\lambda}{b^2}} |a_j| \right) L^b(|a_j|) & \text{if } \lambda > 0, \end{cases} \quad (\text{A.4.2})$$

as $a_j \rightarrow \infty$. The results in Theorem 4 follow by taking derivative respect to $|a_j|$ and b respectively.

A.5 The Calculation of π_j^b

For orthogonal designs, π_j^b with Laplace prior can be integrated out directly from (21)

$$\begin{aligned}\pi_j^b &= \int_0^1 \kappa_j^{\frac{1}{2}} \exp \left[\frac{a_j^2}{2} (1 - \kappa_j) \right] \cdot \frac{1}{2b^2} \kappa_j^{-2} \exp \left(-\frac{1 - \kappa_j}{2b^2 \kappa_j} \right) d\kappa_j \\ &= \frac{1}{2b^2} \left(\frac{2\pi}{\lambda} \right)^{1/2} \exp \left(\frac{1}{2b^2} + \frac{a_j^2}{2} - \frac{\sqrt{a_j^2}}{b} \right) \int_0^1 \left(\frac{\lambda}{2\pi} \right)^{1/2} \kappa_j^{-3/2} \exp \left[-\frac{\lambda(\kappa_j - \mu)^2}{2\mu^2 \kappa_j} \right] d\kappa_j,\end{aligned}\tag{A.5.1}$$

where $\lambda = 1/b^2$ and $\mu = \sqrt{1/(b^2 a_j^2)}$, and the expression in the integral is the CDF of inverse Gaussian distribution. Borrowing the expression of the CDF of the inverse Gaussian, we then integrated out the integral to get expression (26).

π_j^b with horseshoe prior for orthogonal design can also be derived directly from (21):

$$\begin{aligned}\pi_j^b &= \int_0^1 \kappa_j^{\frac{1}{2}} \exp \left[\frac{a_j^2}{2} (1 - \kappa_j) \right] \cdot \frac{b}{\pi} \kappa_j^{-1/2} (1 - \kappa_j)^{-1/2} (1 - \kappa_j + b^2 \kappa_j)^{-1} d\kappa_j \\ &= \frac{1}{\pi b} \exp \left(\frac{a_j^2}{2} \right) \int_0^1 \kappa_j^{1-1} (1 - \kappa_j)^{1/2-1} [(1 - b^{-2})\kappa_j + b^{-2}]^{-1} \exp \left(-\frac{a_j^2}{2} \kappa_j \right) d\kappa_j,\end{aligned}\tag{A.5.2}$$

where the expression in the integral is the transformation of the hypergeometric inverted-beta distribution which was shown to be represented by degenerate hypergeometric functions (Gordy, 1998; Polson and Scott, 2010), thus we can follow Polson and Scott (2010) to express it as (27).

A.6 Lancaster and Šalkauskas Basis for Natural Cubic Spline

In this paper, we follow Chib and Greenberg (2010) to employ the cubic spline LS basis described by Lancaster and Šalkauskas (1986). Consider the j th function $f_j(x)$, and let $\boldsymbol{\nu}_j = (\nu_{1j}, \dots, \nu_{K_j j})$ be the set of $100 \times \frac{k-1}{K_j-1} \%$, $k = 1, \dots, K_j$ quantile of x_{ij} , $i = 1, \dots, n$. Thus $\nu_{1j} = \min_i(x_{ij})$ and $\nu_{K_j j} = \max_i(x_{ij})$. K_j denotes the number of knots for the spline

functions. Then the cubic spline expansion of $f_j(x)$ is expressed as

$$\begin{aligned} f_j(x_{ij}) &= \sum_{k=1}^{K_j} [\Phi_{kj}(x_{ij})g_{kj} + \Psi_{kj}(x_{ij})s_{kj}] \\ &= \Phi_j(x_{ij})^T \mathbf{g}_j + \Psi_j(x_{ij})^T \mathbf{s}_j, \end{aligned} \quad (\text{A.6.1})$$

where $\mathbf{g}_j = (g_{1j}, \dots, g_{K_j j})^T$ and $\mathbf{s}_j = (s_{1j}, \dots, s_{K_j j})^T$ are the coefficients of this expression, $\Phi_j(x_{ij}) = [\Phi_{1j}(x_{ij}), \dots, \Phi_{K_j j}(x_{ij})]^T$ and $\Psi_j(x_{ij}) = [\Psi_{1j}(x_{ij}), \dots, \Psi_{K_j j}(x_{ij})]^T$ are two basis vectors, and the basis functions $\{\Phi_{kj}(x)\}_{k=1}^{K_j}$ and $\{\Psi_{kj}(x)\}_{k=1}^{K_j}$ are defined as

$$\begin{aligned} \Phi_{kj}(x) &\propto \begin{cases} 0, & x < \nu_{k-1,j} \\ -(2/h_{kj}^3)(x - \nu_{k-1,j})^2(x - \nu_{kj} - 0.5h_{kj}), & \nu_{k-1,j} \leq x < \nu_{kj} \\ (2/h_{k+1,j}^3)(x - \nu_{k+1,j})^2(x - \nu_{kj} + 0.5h_{k+1,j}), & \nu_{kj} \leq x < \nu_{k+1,j} \\ 0 & x \geq \nu_{k+1,j}, \end{cases} \\ \Psi_{kj}(x) &\propto \begin{cases} 0, & x < \nu_{k-1,j} \\ (1/h_{kj}^2)(x - \nu_{k-1,j})^2(x - \nu_{kj}), & \nu_{k-1,j} \leq x < \nu_{kj} \\ (1/h_{k+1,j}^2)(x - \nu_{k+1,j})^2(x - \nu_{kj}), & \nu_{kj} \leq x < \nu_{k+1,j} \\ 0 & x \geq \nu_{k+1,j}, \end{cases} \end{aligned} \quad (\text{A.6.2})$$

where $h_{kj} = \nu_{kj} - \nu_{k-1,j}$. Note that Φ_{1j}, Ψ_{1j} and $\Phi_{K_j j}, \Psi_{K_j j}$ are defined by last two lines and first two lines of above expressions respectively. \mathbf{g}_j and \mathbf{s}_j are interpreted the ordinate and slope of $f_j(x)$. Since $f_j(x)$ is a natural cubic splines with the second derivative equal to zero at two end points, and continuous derivative at knot points, both \mathbf{g}_j and \mathbf{s}_j are constrained (Lancaster and Šalkauskas, 1986) by $\mathbf{s}_j = A_j^{-1} C_j \mathbf{g}_j$, where

$$A_j = \begin{pmatrix} 2 & 1 & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\ \omega_{2j} & 2 & \mu_{2j} & 0 & 0 & \cdots & 0 & 0 & 0 \\ 0 & \omega_{2j} & 2 & \mu_{2j} & 0 & \cdots & 0 & 0 & 0 \\ \vdots & \cdots & \ddots & \ddots & \ddots & \cdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 & \cdots & \omega_{K_j-1,j} & 2 & \mu_{K_j-1,j} \\ 0 & 0 & 0 & 0 & 0 & \cdots & 0 & 1 & 2 \end{pmatrix},$$

and

$$C_j = \begin{pmatrix} -\frac{1}{h_{2j}} & \frac{1}{h_{2j}} & 0 & 0 & \cdots & 0 & 0 & 0 \\ -\frac{\omega_{2j}}{h_{2j}} & \frac{\omega_{2j}}{h_{2j}} - \frac{\mu_{2j}}{h_{3j}} & \frac{\mu_{2j}}{h_{3j}} & 0 & \cdots & 0 & 0 & 0 \\ 0 & -\frac{\omega_{3j}}{h_{3j}} & \frac{\omega_{3j}}{h_{3j}} - \frac{\mu_{3j}}{h_{4j}} & \frac{\mu_{3j}}{h_{4j}} & \cdots & 0 & 0 & 0 \\ \vdots & \cdots & \ddots & \ddots & \cdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & -\frac{\omega_{K_j-1,j}}{h_{K_j-1}} & \frac{\omega_{K_j-1,j}}{h_{K_j-1}} - \frac{\mu_{K_j-1,j}}{h_{K_j}} & \frac{\mu_{K_j-1,j}}{h_{K_j}} \\ 0 & 0 & 0 & 0 & \cdots & 0 & -\frac{1}{h_{K_jj}} & \frac{1}{h_{K_jj}} \end{pmatrix},$$

where $\omega_{kj} = h_{kj}/(h_{kj} + h_{k+1,j})$ and $\mu_{kj} = 1 - \omega_{kj}$ for $k = 2, \dots, K_j$. With this constraints, \mathbf{s}_j can be replaced from the function expression (A.6.1),

$$\begin{aligned} f_j(x_{ij}) &= [\Phi_j(x_{ij})^T + \Psi_j(x_{ij})^T A_j^{-1} C_j] \mathbf{g}_j \\ &= \mathbf{t}_j^T(x_{ij}) \mathbf{g}_j, \end{aligned} \tag{A.6.3}$$

where $\mathbf{t}_j^T(x_{ij}) = (t_{1j}(x_{ij}), \dots, t_{K_jj}(x_{ij})) = \Phi_j(x_{ij})^T + \Psi_j(x_{ij})^T A_j^{-1} C_j$. Furthermore, consider the identifying constraints, $\sum_k g_{kj} = 0$, we can express $g_{1j} = -(g_{2j} + \dots + g_{K_jj})$, thus

$$f_j(x_{ij}) = \mathbf{t}_j^T(x_{ij}) \mathbf{g}_j = [t_{2j}(x_{ij}) - t_{1j}(x_{ij})]g_{2j} + \dots + [t_{K_jj}(x_{ij}) - t_{1j}(x_{ij})]g_{K_jj} = \mathbf{z}_j^{*T}(x_{ij}) \boldsymbol{\beta}_j,$$

where $\boldsymbol{\beta}_j = (g_{2j}, \dots, g_{K_jj})^T$ and we define matrix

$$Z_j^* = \begin{pmatrix} \mathbf{z}_j^{*T}(x_{1j}) \\ \vdots \\ \mathbf{z}_j^{*T}(x_{n_jj}) \end{pmatrix}.$$

Now the j th nonparametric function expressed by the natural cubic spline basis is $f_j(\mathbf{x}_j) = Z_j^* \boldsymbol{\beta}_j$. In order to incorporate the assumption of a priori smoothness, Chib and Greenberg (2010) consider a prior distribution on $\boldsymbol{\beta}_j$'s as,

$$[\boldsymbol{\beta}_j | \sigma_{ej}^2, \sigma_{dj}^2] \sim N[\mathbf{0}, \Delta_j^{-1} T_j (\Delta_j^{-1})^T], \tag{A.6.4}$$

where N is the $K_j - 1$ dimensional multivariate normal distribution, and

$$T_j = \begin{pmatrix} \sigma_{ej}^2 & 0 & 0 \\ 0 & \sigma_{dj}^2 I_{K_j-3} & 0 \\ 0 & 0 & \sigma_{ej}^2 \end{pmatrix},$$

where two variance components σ_{ej}^2 and σ_{dj}^2 are selected here because of the different normal assumptions for the differences of the ordinates and the differences of slopes. Δ_j is given by

$$\Delta_j = \begin{pmatrix} \frac{2}{h_{2j}} & \frac{1}{h_{2j}} & \frac{1}{h_{2j}} & \frac{1}{h_{2j}} & \frac{1}{h_{2j}} & \dots & \frac{1}{h_{2j}} & \frac{1}{h_{2j}} \\ \frac{1}{h_{2j}} & -\left(\frac{1}{h_{2j}} + \frac{1}{h_{3j}}\right) & \frac{1}{h_{3j}} & 0 & 0 & \dots & 0 & 0 \\ 0 & \frac{1}{h_{3j}} & -\left(\frac{1}{h_{3j}} + \frac{1}{h_{4j}}\right) & \frac{1}{h_{4j}} & 0 & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \dots & \vdots & \vdots \\ 0 & 0 & 0 & 0 & 0 & \dots & -\frac{1}{h_{K_jj}} & \frac{1}{h_{M_jj}} \end{pmatrix}.$$

So far the construction of function $f_j(\mathbf{x}_j)$ is exactly the same as Chib and Greenberg (2010). Note that $f_j(\mathbf{x}_j) = Z_j^* \boldsymbol{\beta}_j$ with the prior of $\boldsymbol{\beta}_j$ given by (A.6.4) is equivalent to have $f_j(\mathbf{x}_j) = Z_j^* \Delta_j^{-1} \boldsymbol{\beta}_j$ with $[\boldsymbol{\beta}_j] \sim N(\mathbf{0}, T_j)$. Henceforth, we define the final $n \times M_j$ basis matrix $Z_j = Z_j^* \Delta_j^{-1}$ such that $f_j(\mathbf{x}_j) = Z_j \boldsymbol{\beta}_j$, where $M_j = K_j - 1$. Define $\tau_{ej} = \sigma_{ej}^{-2}$ and $\tau_{bj} = \sigma_{bj}^{-2}$, and modify the one variance component prior algorithm in Section 7, then we can easily employ the LS basis into BSAM.

References

- Bae, K. and Mallick, B. (2004). Gene Selection Using a Two-level Hierarchical Bayesian Model. *Bioinformatics*, **20**, 3423-3430.
- Barndorff-Nielsen, O., Kent, J., and Sørensen, M. (1982). Normal Variance-Mean Mixtures and z Distribution. *International Statistical Review*, **50**, 145-159.
- Carlin, B. P. and Polson, N. G. (1991). Inference for Nonconjugate Bayesian Models Using the Gibbs Sampler. *The Canadian Journal of Statistics*, **19**, 399-405.
- Carvalho, C. and Polson, N. G. (2010). The Horseshoe Estimator for Sparse Signals. *Biometrika*, **97**, 465-480.
- Chib, S. and Greenberg, E. (2010). Additive Cubic Spline Regression with Dirichlet Process Mixture Errors. *Journal of Econometrics*, **156**, 322-336.
- Gelman, A. (2006). Prior Distribution for Variance Parameters in Hierarchical Models. *Bayesian Analysis*, **1**, 515-533.
- Geyer, C. J. and Thompson, E. A. (1995). Annealing Markov Chain Monte Carlo with Application to Ancestral Inference. *Journal of the American Statistical Association*, **90**, 909-920.
- Gordy, M. B. (1998). A generalization of Generalized Beta Distribution. In *Finance and Economics Discussion Series*. Board of Governors of the Federal Reserve System.
- George, E. I. and McCulloch, R. R. (1993). Variable Selection via Gibbs Sampling. *Journal of the American Statistical Association*, **88**, 881-889.
- Hans, C. M. (2009). Bayesian Lasso Regression. *Biometrika*, **96**, 835-845.

- Heaton, M. and Scott, J. (2010). Bayesian Computation and the Linear Model. In *Frontiers of Statistical Decision Making and Bayesian Analysis*, eds. M. H. Chen, D. K. Dey, P. Mueller, D. Sun, and K. Ye. New York: Springer.
- Higdon, D. M. (1998). Auxiliary Variable Methods for Markov Chain Monte Carlo with Applications. *Journal of the American Statistical Association*, **93**, 585-595.
- Iba, Y. (2001). Extended Ensemble Monte Carlo. *International Journal of Modern Physics C*, **12**, 623-656.
- Ishwaran, H. I. and Rao, J. S. (2005). *The Annals of Statistics*, **33**, 730-773.
- Kuo, L. and Mallick, B. (1998). Variable Selection for Regression Models. *Sankhyā: The Indian Journal of Statistics B*, **60**, 65-81.
- Lancaster, P. and Šalkauskas, K. (1986). *Curve and Surface Fitting: an Introduction*. San Diego: Academic Press.
- Li, F. and Zhang, N. R. (2010). Bayesian Variable Selection in Structured High-Dimensional Covariates Spaces with Applications in Genomics. *Journal of the American Statistical Association*, **105**, 1202-1214.
- Liang, F., Paulo, R., Molina, G., Clyde, M. A., and Berger, J. O. (2008). Mixtures of g Priors for Bayesian Variable Selection. *Journal of the American Statistical Association*, **103**, 410-423.
- Lin, Y. and Zhang, H. H. (2006). Component Selection and Smoothing in Multivariate Nonparametric Regression. *The Annals of Statistics*, **34**, 2272-2297.
- Lykou, A. and Ntzoufras, I. (2012). On Bayesian Lasso Variable Selection and the Specifi-

- cation of the Shrinkage Parameter. *Statistics and Computing*, DOI: 10.1007/s11222-012-9316-x, 2012.
- Lyubartsev, A. P., Martsinovski, A. A., Shevkunov, S. V., and Vorontsov-Velyaminov, P. N. (1992). New Approach to Monte Carlo Calculation of the Free Energy: Method of Expanded Ensembles. *Journal of Chemical Physics*, **96**, 1776-1783.
- Park, T. and Casella, G. (2008). The Bayesian Lasso. *Journal of the American Statistical Association*, **103**, 681-686.
- Monni, S. and Li H. (2010). Bayesian Methods for Network-Structured Genomics Data. *UPenn Biostatistics Working Papers*, Working Paper 34.
- Mootha, V. K., Lindgren, C. M., Eriksson, K., Subramanian, A., Sihag, S., Lehar, J., Puigserver, P., Carlsson, E., Ridderstrale, M., Laurila, E., Houstis, N., Daly, M. J., Patterson, N., Mesirov, J. P., Golub, T. R., Tamayo, P., Spiegelman, B., Lander, E. S., Hirschhorn, J. N., Altshuler, D., and Groop, L. C. (2003). PGC-1 alpha-Responsive Genes Involved in Oxidative Phosphorylation are Coordinately Downregulated in Human Diabetes. *Nature Genetics*, **34**, 267-273.
- Newman, M. E. J. and Barkema, G. T. (1999). *Monte Carlo Methods in Statistical Physics*. New York: Oxford University Press.
- Nott, D. J. and Green, P. J. (2004). Bayesian Variable Selection and the Swenden-Wang Algorithm. *Journal of Computational and Graphical Statistics*, **13**, 141-157.
- Pang, H., Lin, A., Holford, M., Enerson, B., Lu, B., Lawton, M. P., Floyd, E., and Zhao, H. (2006). Pathway Analysis Using Random Forests Classification and Regression. *Bioinformatics*, **22**, 2028-2036.

- Pericchi, L. R. and Smith, A. (1992). Exact and Approximate Posterior Moments for a Normal Location Parameter. *Journal of the Royal Statistical Society, Series B*, **54**, 793-804
- Polson, N. G. and Scott, J. G. (2010). On the Half-Cauchy Prior for a Global Scale Parameter. *Technical report*, University of Texas at Austin.
- Polson, N. G. and Scott, J. G. (2011). Shrink Globally, Act Locally: Sparse Bayesian Regularization and Prediction. In *Proceedings of the 9th Valencia World Meeting on Bayesian Statistics*. Oxford University Press, 501-538.
- Polson, N. G. and Scott, J. G. (2012). Local Shrinkage Rules, Lévy Processes and Regularization Regression. *Journal of the Royal Statistical Society, Series B*, **74**, 287-311.
- Reich, B. J., Storlie, C. B., and Bondell, H. D. (2009). Variable Selection in Bayesian Smoothing Spline ANOVA Models: Application to Deterministic Computer Codes. *Journal of Econometrics*, **51**, 110-119.
- Scheipl, F. (2011). spikeSlabGAM: Bayesian Variable Selection, Model Choice and Regularization for Generalized Additive Mixed Models in R. *Journal of Statistical Software*, **43**, 1-24.
- Smith, M. and Kohn, R. (1996). Nonparametric Regression Using Bayesian Variable Selection. *Journal of Econometrics*, **75**, 317-343.
- Stingo, F. C., Chen, Y. A., Tadesse, M. G., and Vannucci, M. (2011). Incorporating Biological Information into Linear Models: a Bayesian Approach to the Selection of Pathways and Genes. *The Annals of Applied Statistics*, **5**, 1978-2002.
- Swendsen, R. H. and Wang, J. S. (1987). Nonuniversal critical dynamics in Monte Carlo simulations. *Physical Review Letters*, **58**, 86-88.

- Tai, F., Pan, W., and Shen, X. (2010). Bayesian Variable Selection in Regression with Networked Predictors. In *High-Dimensional Data Analysis*, eds. T. Cai and X. Shen. Singapore: World Scientific, 147-165.
- Tipping, M. (2001). Sparse Bayesian Learning and the Relevance Vector Machine. *Journal of Machine Learning Research*, **1**, 211-244.
- West, M. (1987). On Scale Mixtures of Normal Distributions. *Biometrika*, **74**, 646-648.
- Wolff, U. (1989). Collective Monte Carlo Updating for Spin Systems. *Physical Review Letters*, **62**, 361-364.